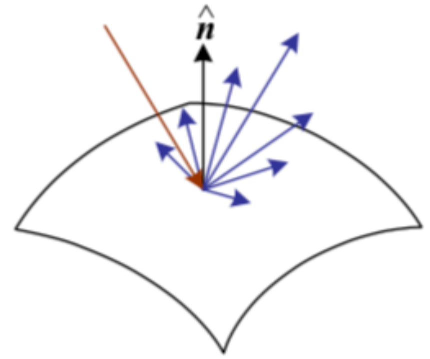

2. Image Formation

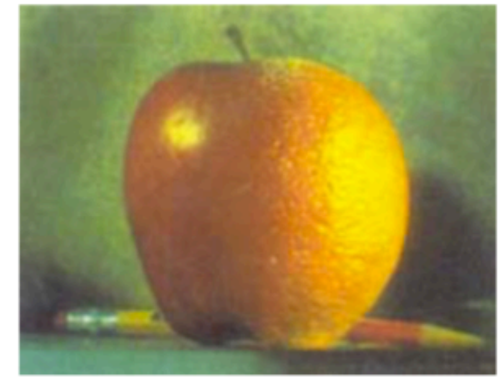

3. Image Processing

4. Features

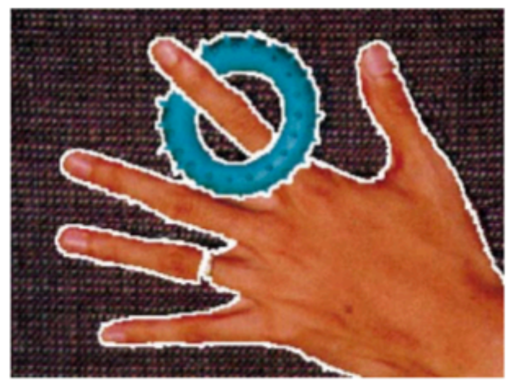

5. Segmentation

6-7. Structure from Motion

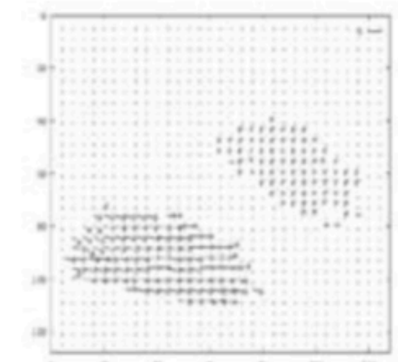

8. Motion

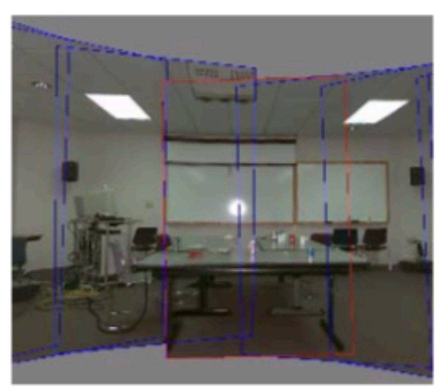

9. Stitching

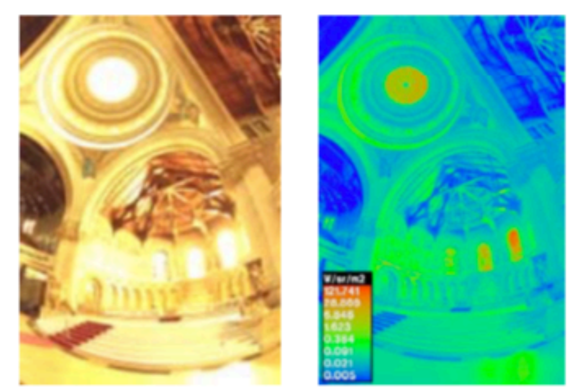

10. Computational Photography

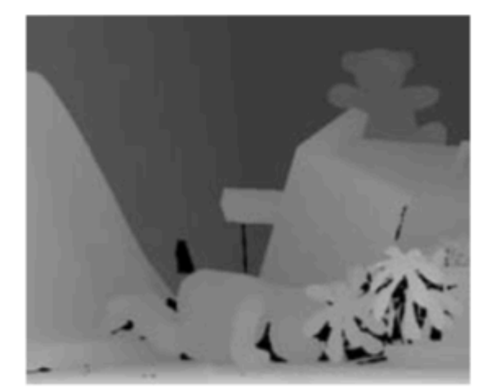

11. Stereo

12. 3D Shape

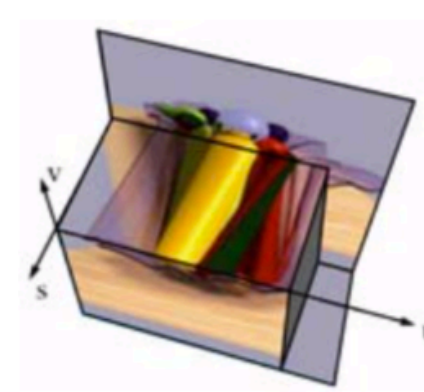

13. Image-based Rendering

14. Recognition

Stitching and HDR,

(Lords Cricket Ground, London, UK, by I. Essa)

Slides by Irfan Essa

Adapted for CS 4476 by Frank Dellaert

More detailsin Szeliski Ch. 9, Ch 10.

Georgia Tech College of Computing

5 Steps to Make a Panorama

(Lords Cricket Ground, London, UK, by I. Essa)

* ✓Capture Images
* ✓Detection and matching
* ✓Warping → Aligning Images
* ✓Blending, Fading, Cutting
* Cropping (Optional)

Align Images: Translate??

L

R

L on top

R on top

Better: Warp

A Bundle of Rays Contains all Views

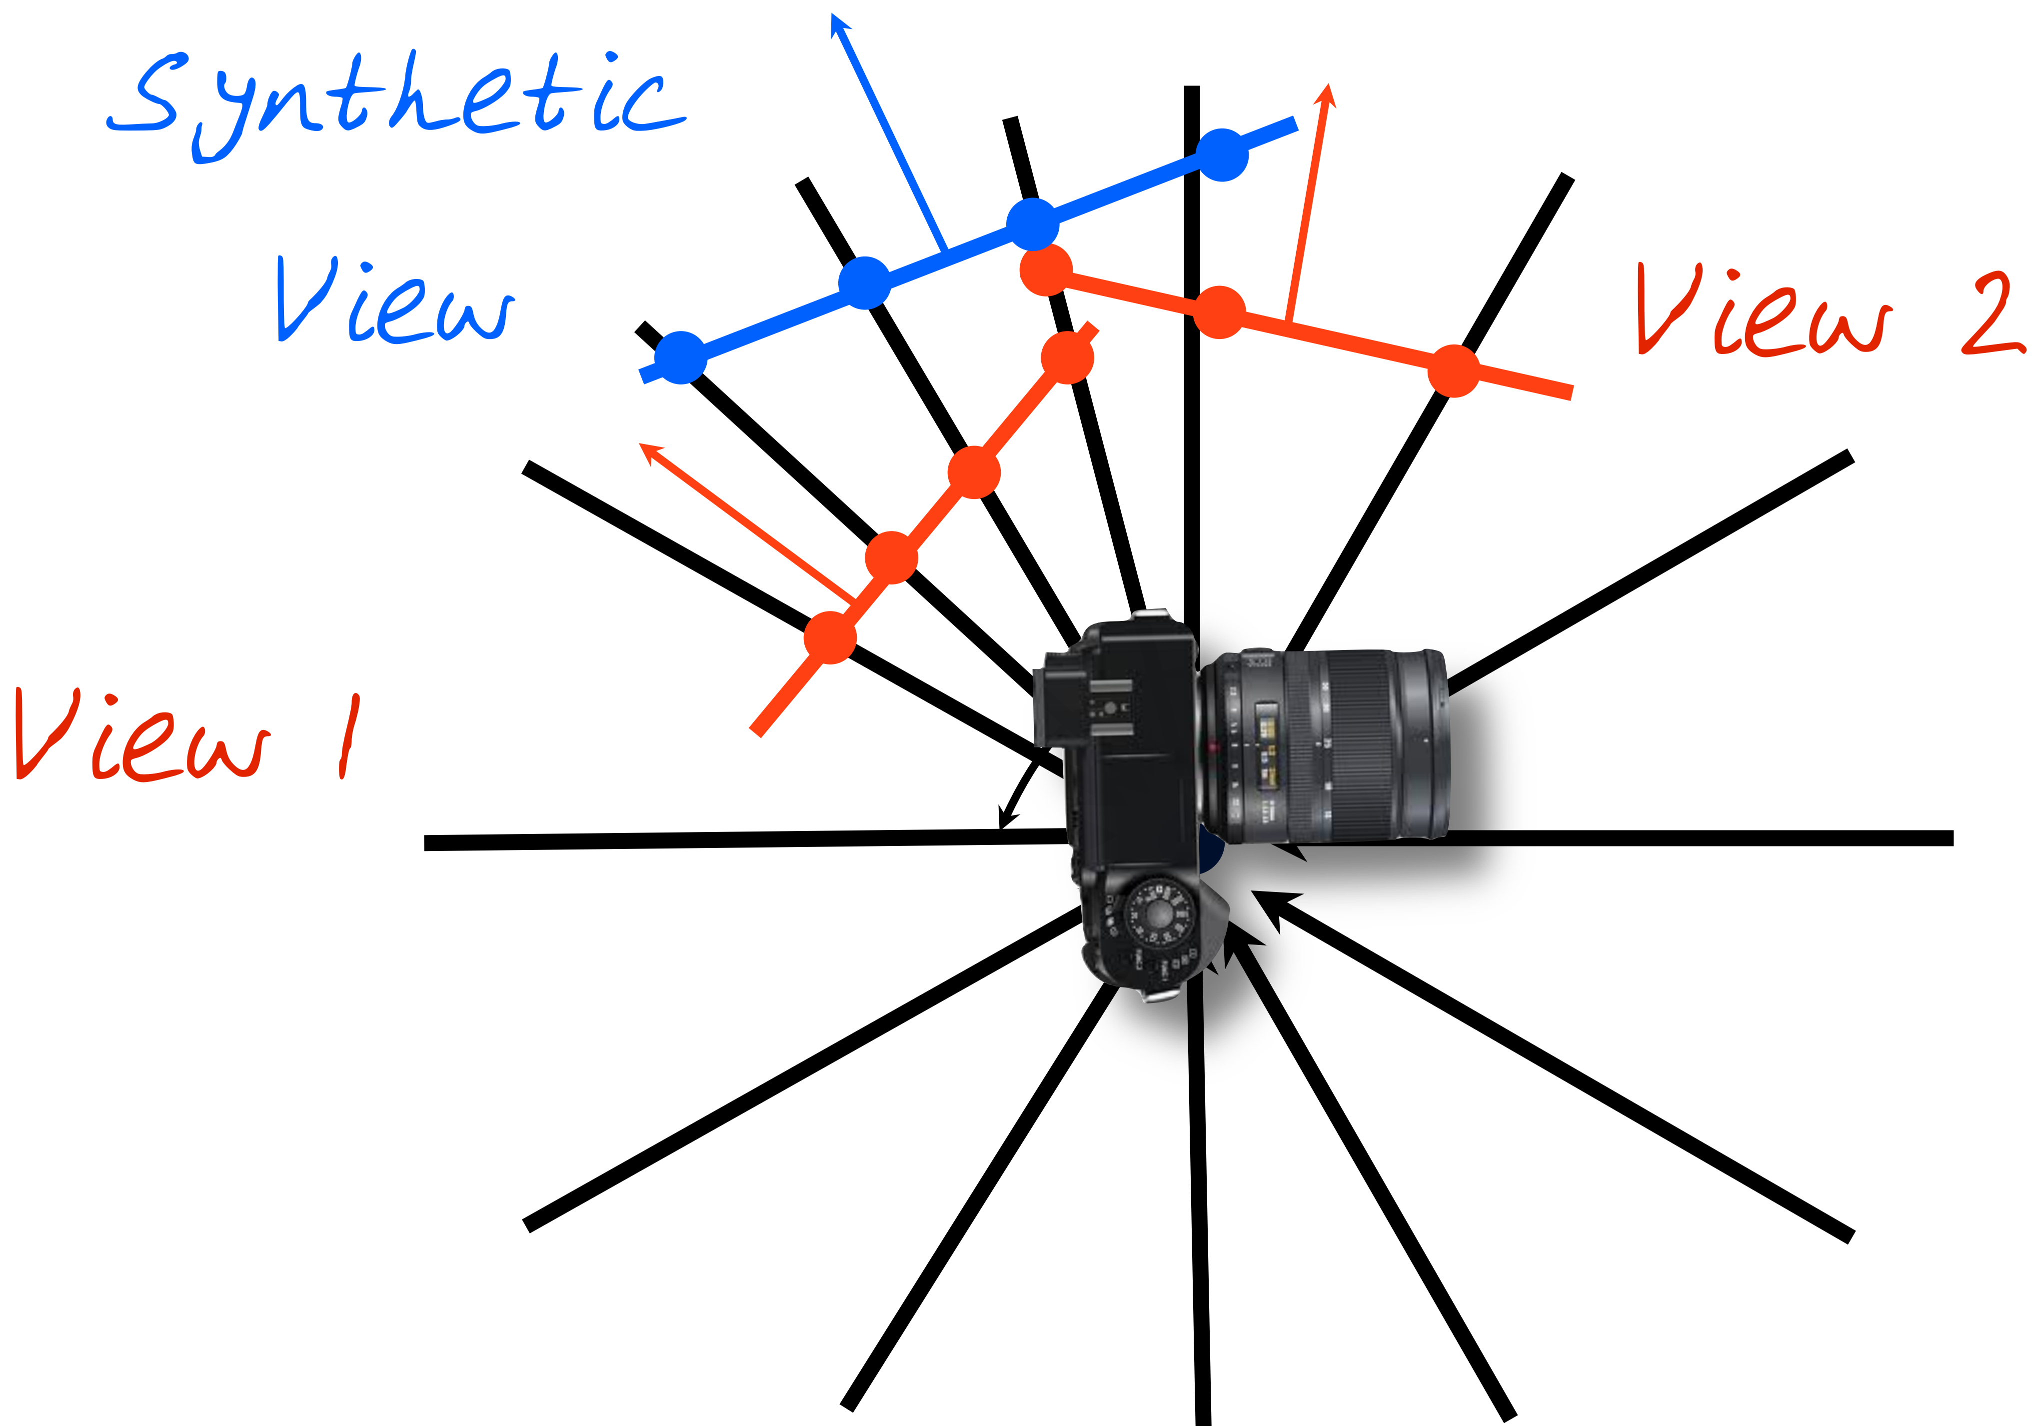

Possible to generate any synthetic camera view as long as it has the same center of projection!

Slide motivated by James Hays

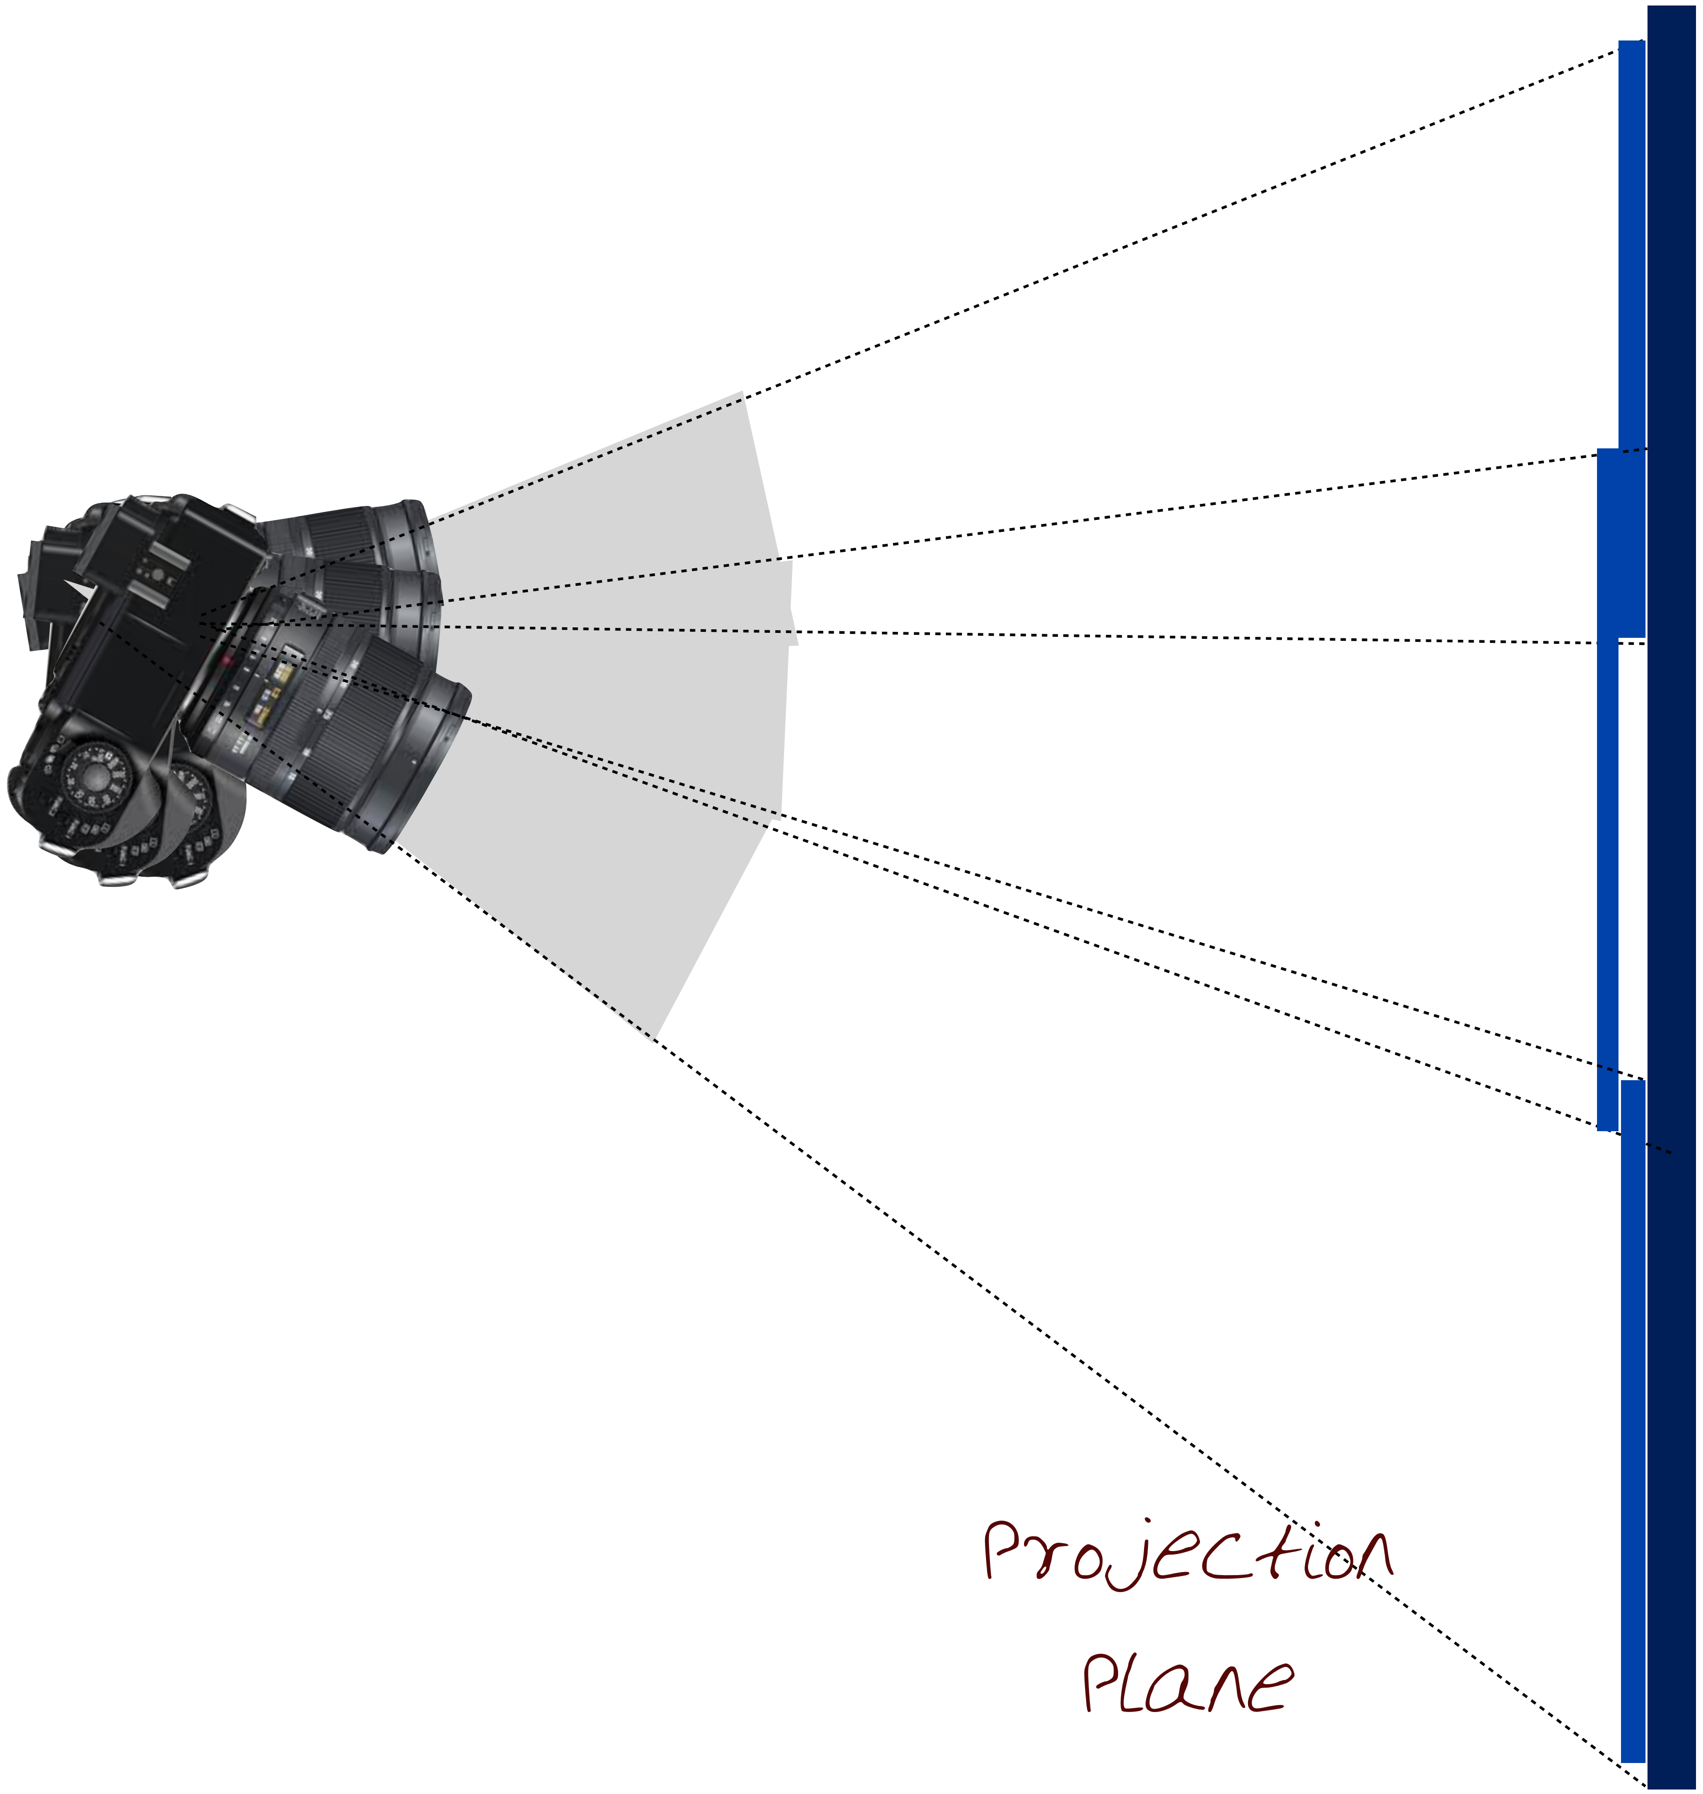

Image Re-projection to Panorama Projection Plane

* The panorama mosaic has a natural interpretation in 3D
  * Images are reprojected onto a common plane
  * The mosaic is formed on this plane
  * Mosaic is a synthetic wide-angle camera

Image Re-Projection (1)

To relate two images from the same camera center and map a pixel from PP1 to PP2:

- Cast a ray through each pixel in PP1
- Draw the pixel where that ray intersects PP2

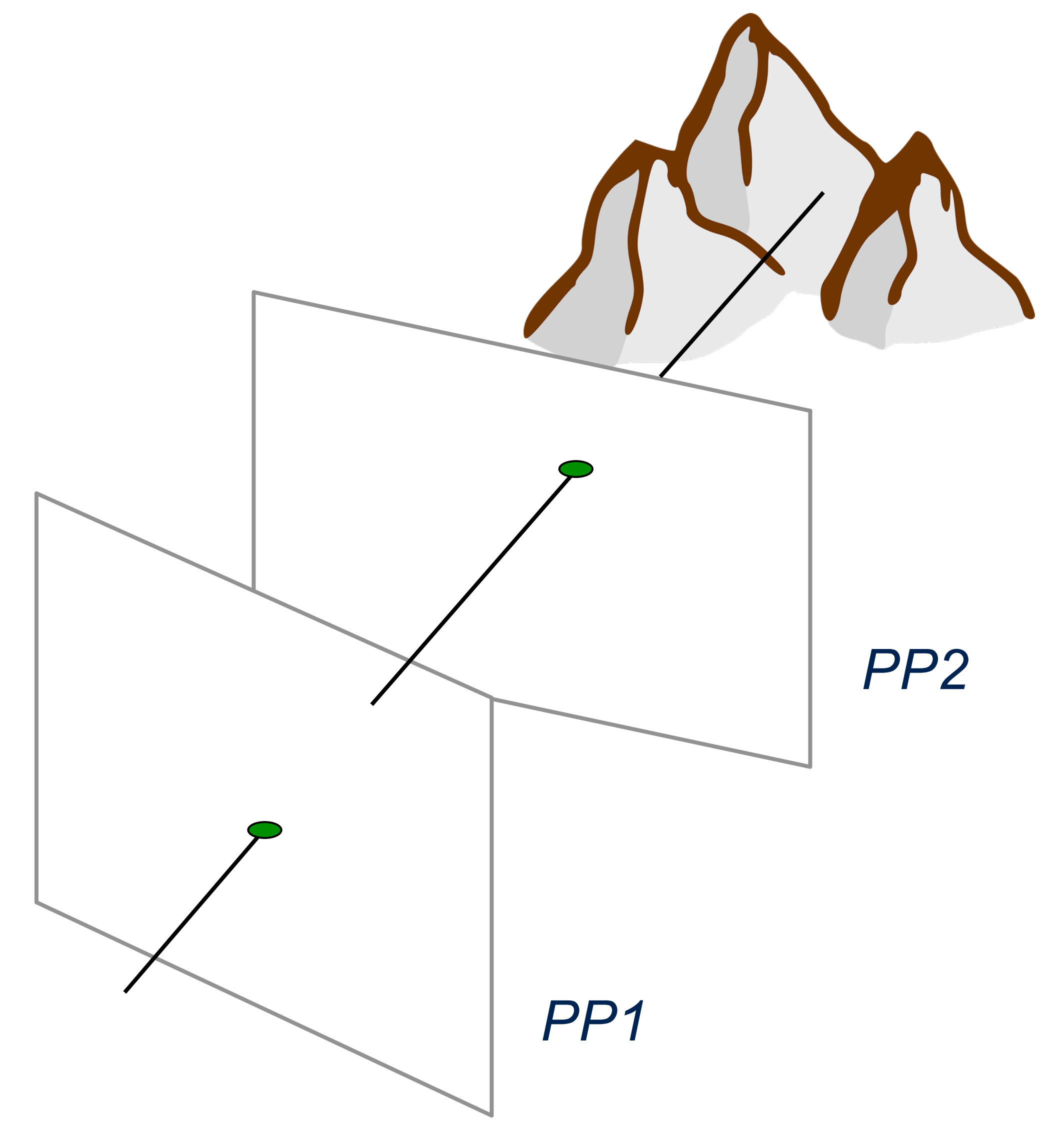

Image Re-Projection (2)

To relate two images from the same camera center and map a pixel from PP1 to PP2:

* Rather than a 3D re-projection,
* Think of it as a 2D image warp from one image to another
* Do not need to know the geometry of the two planes with respect to the eye?

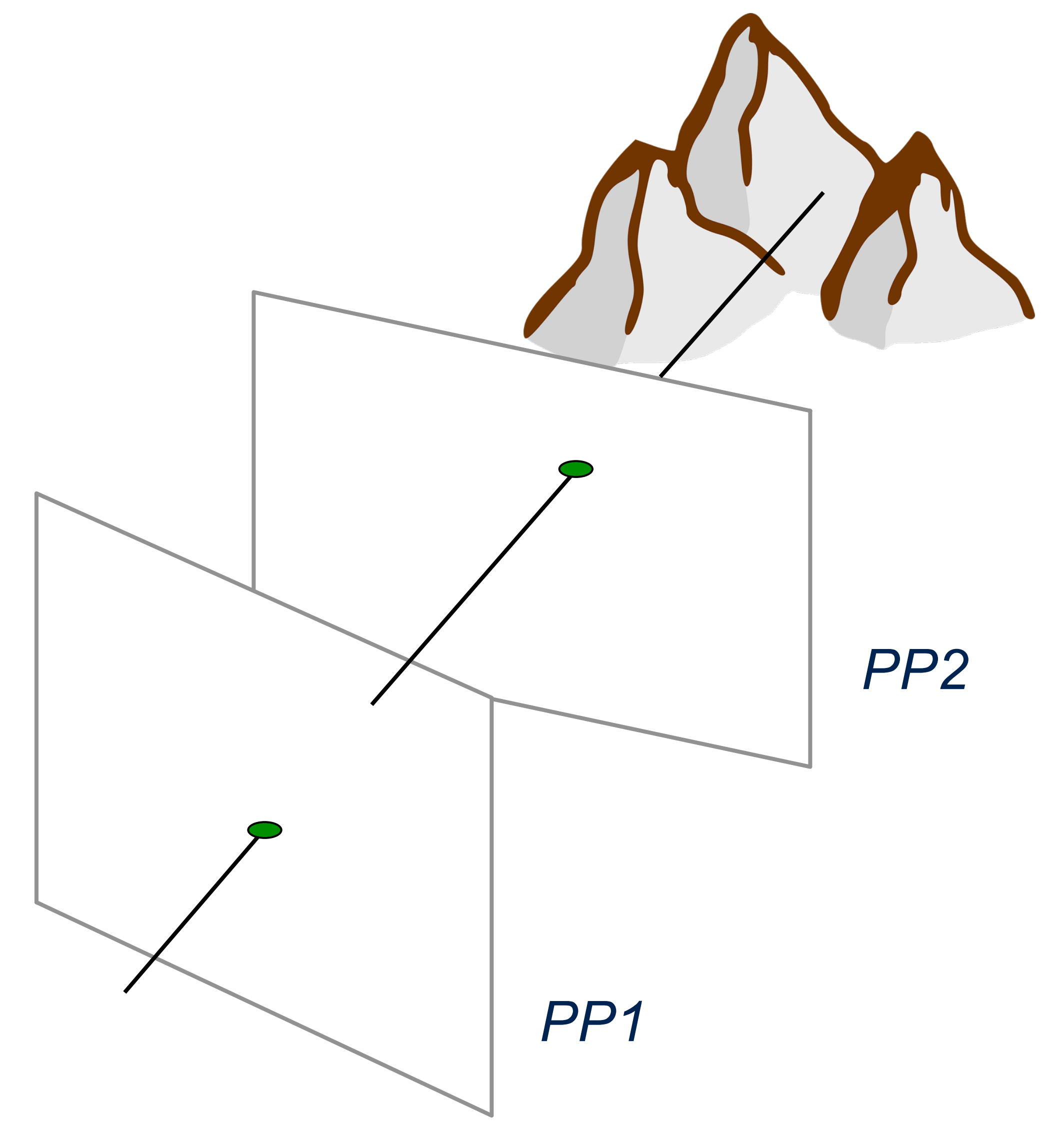

Recall: Image Transforms

Which transform is the right one for warping **PP1** into **PP2**?

E.g. translation, Euclidean, affine, projective

Translation: 2 unknowns, Euclidean: 3 unknowns

Affine: 6 unknowns, Projective: 8 unknowns

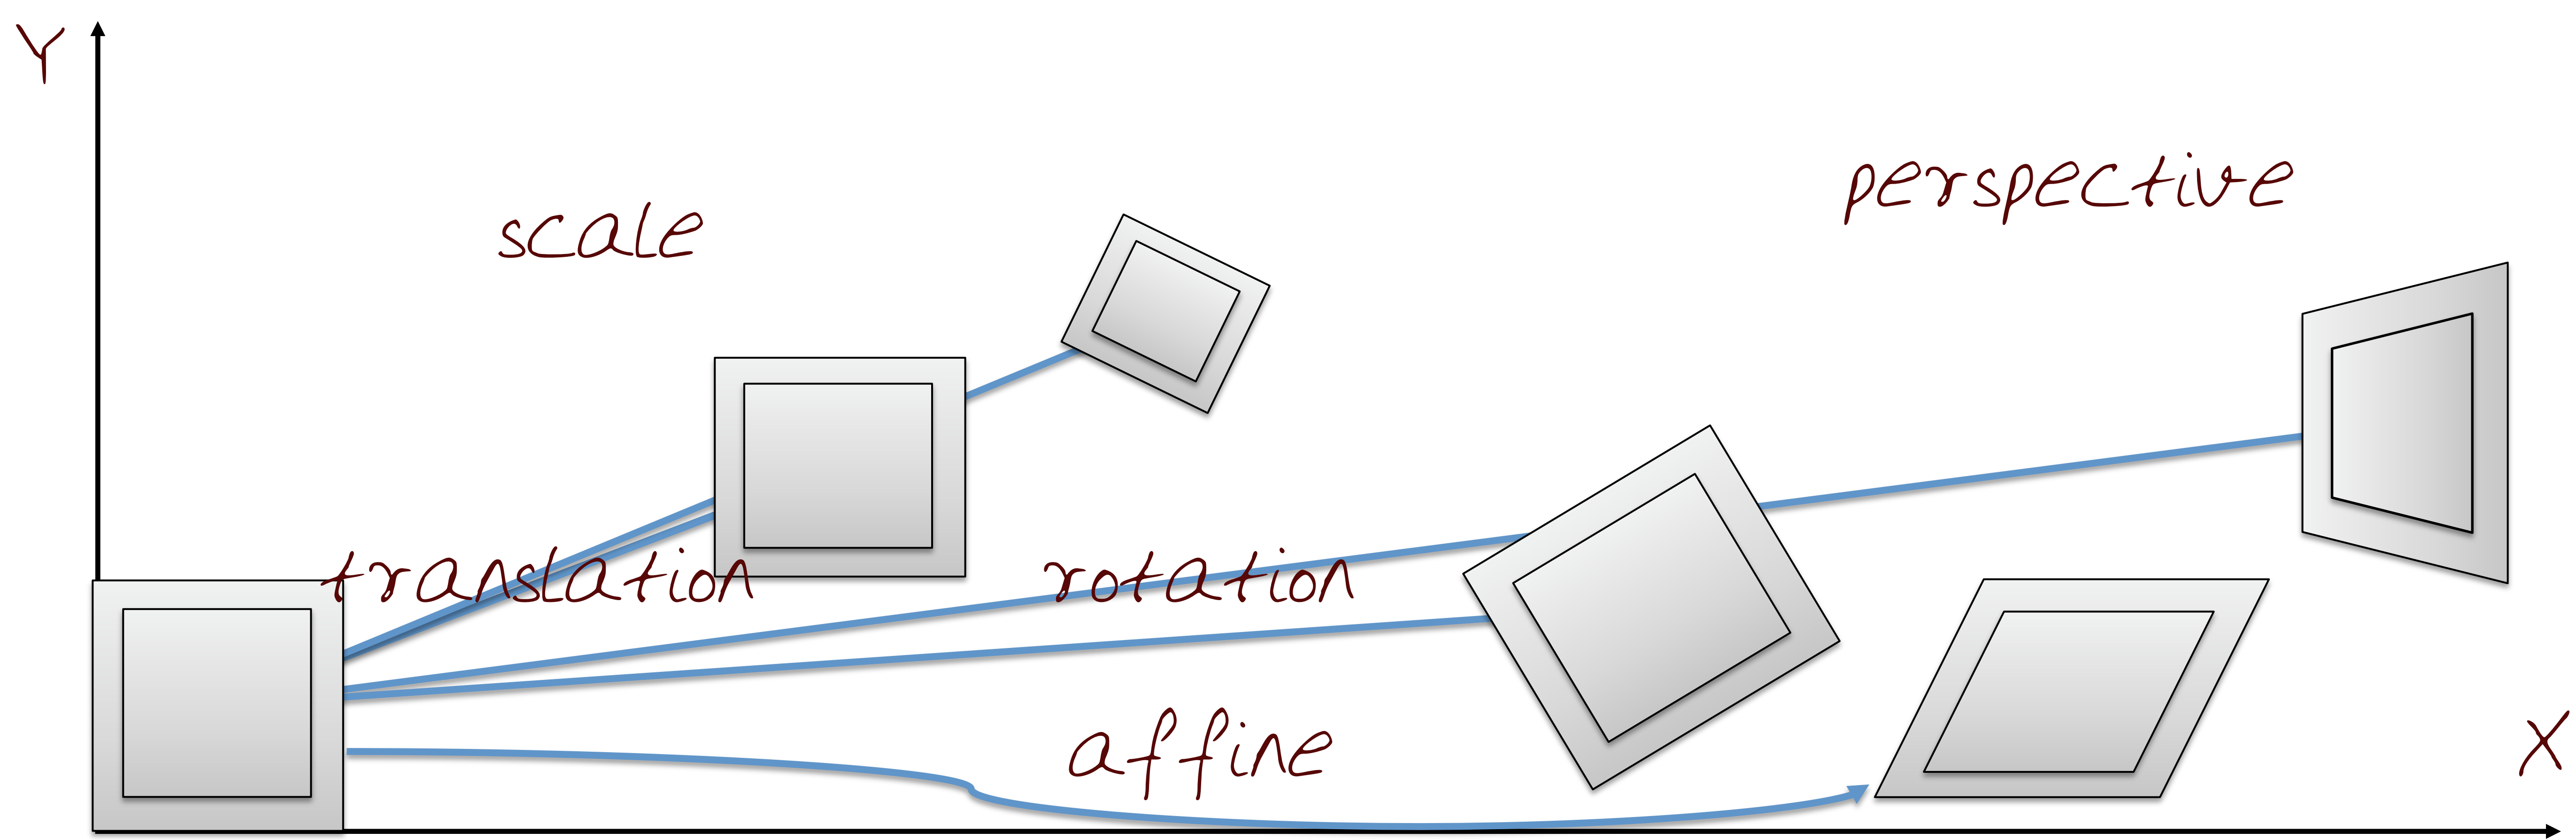

Homography Recap

Relates two images from the same camera center

* Rectangle should map to arbitrary quadrilateral
* Parallel lines aren't parallel
* Straight lines must be straight

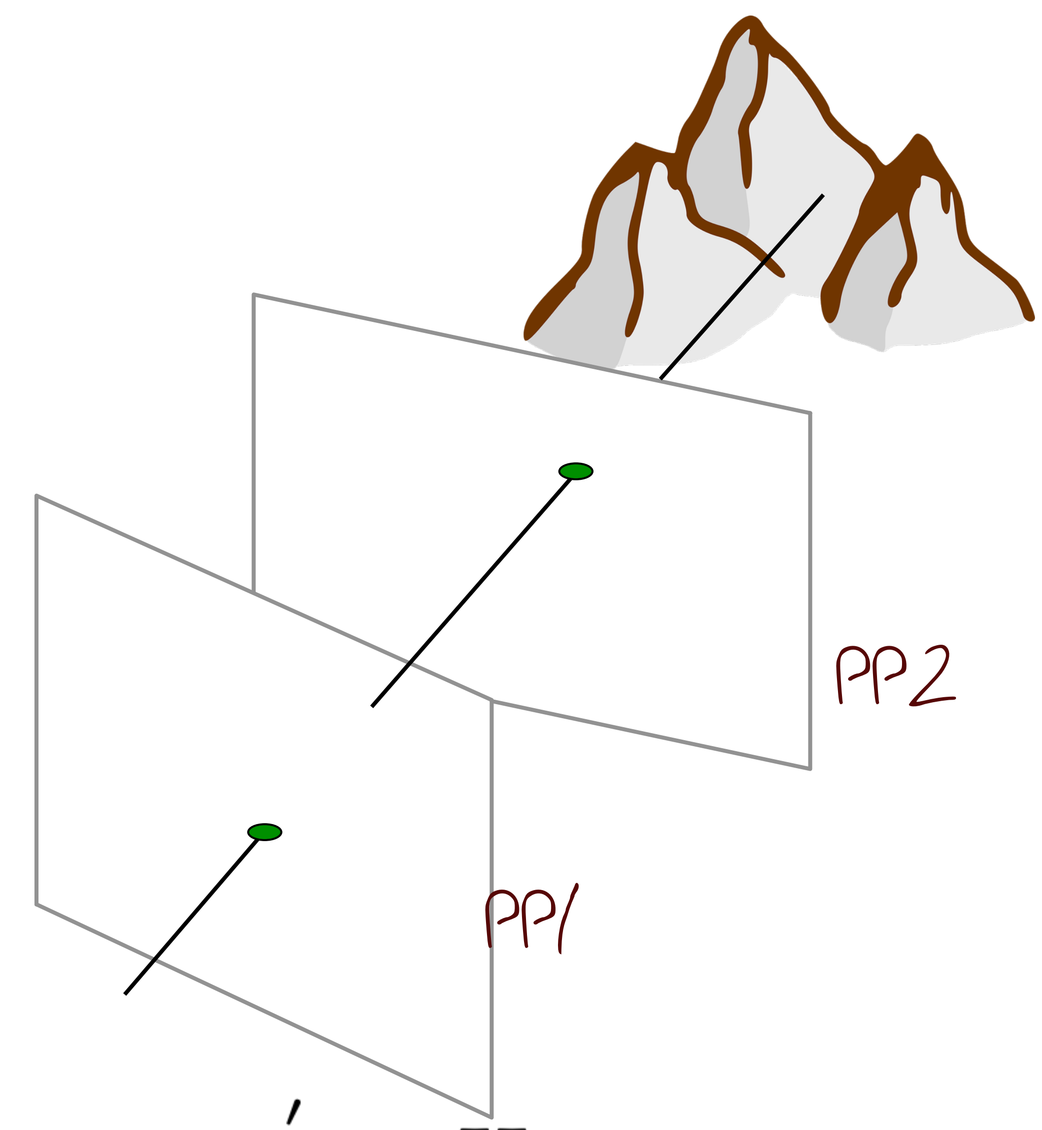

$$p' = Hp$$

$$\begin{bmatrix} wx' \\ wy' \\ w \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

http://en.wikipedia.org/wiki/Homography

Computing Homographies
Nonlinear Least-Squares (Ch. 6)

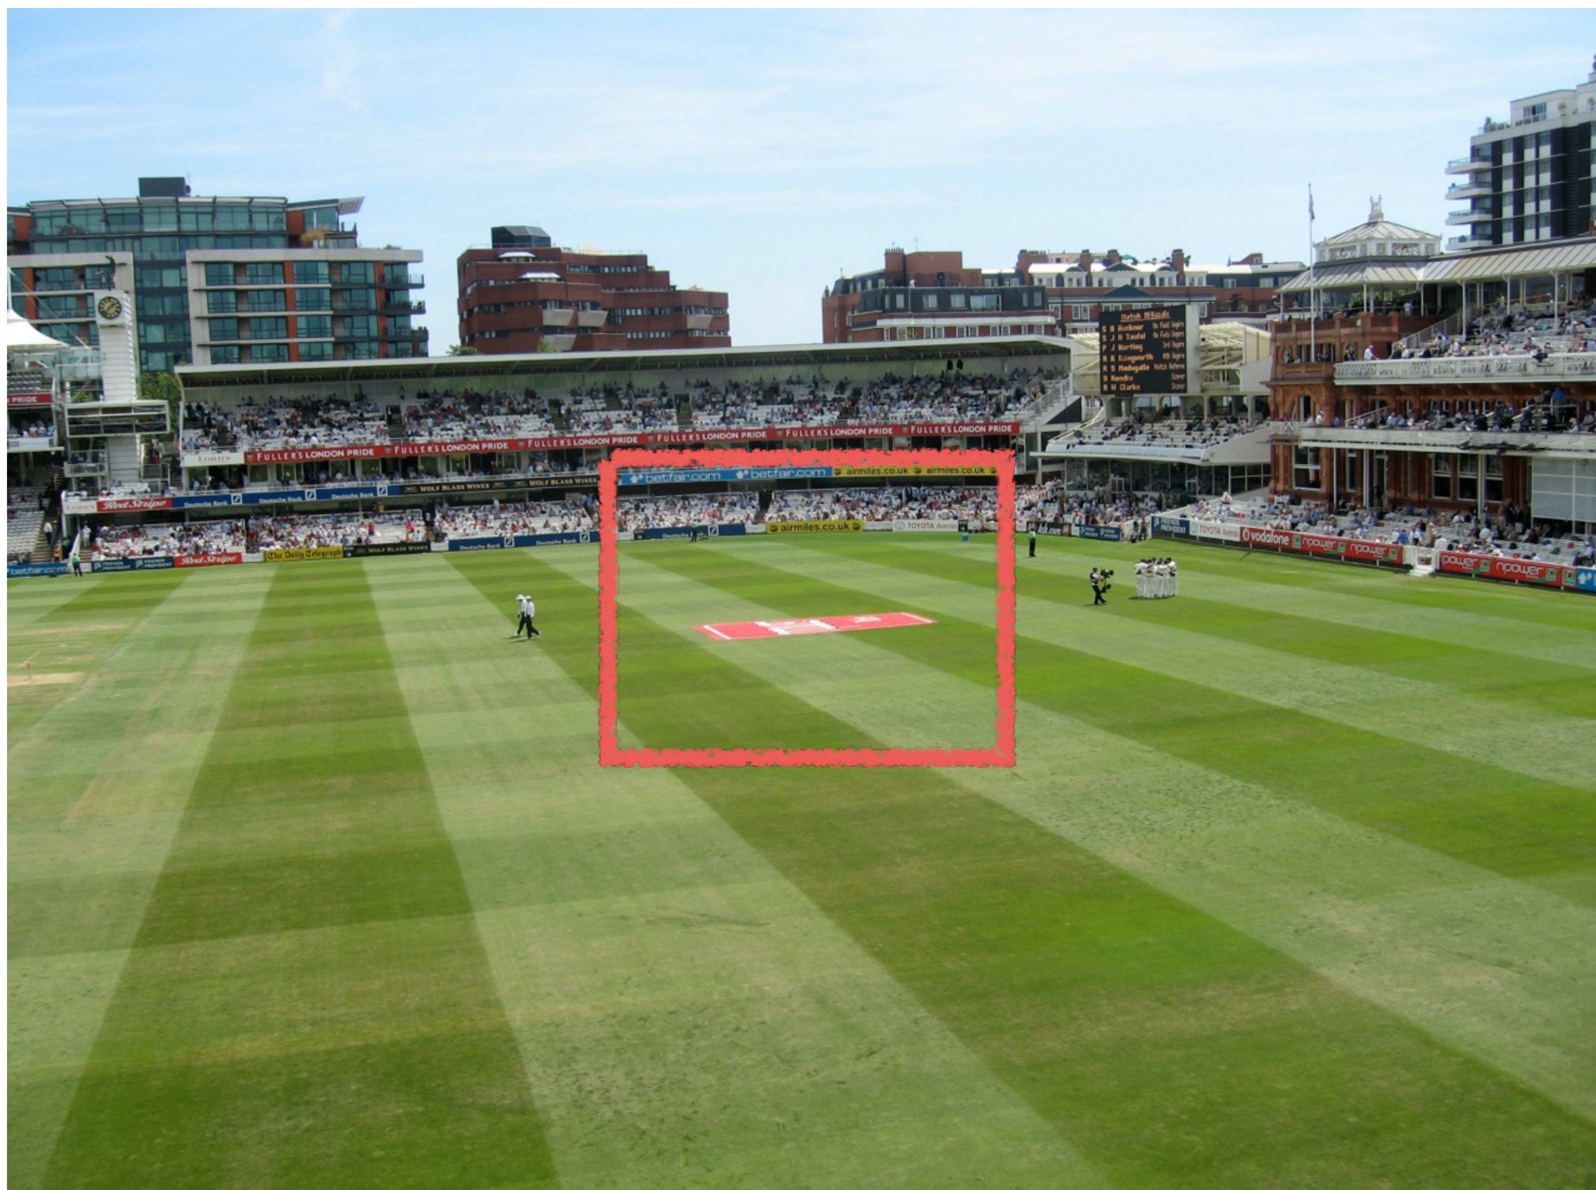

Warp into a Shared Coordinate Space

Warping and Interpolation

Dealing with BAD Matches

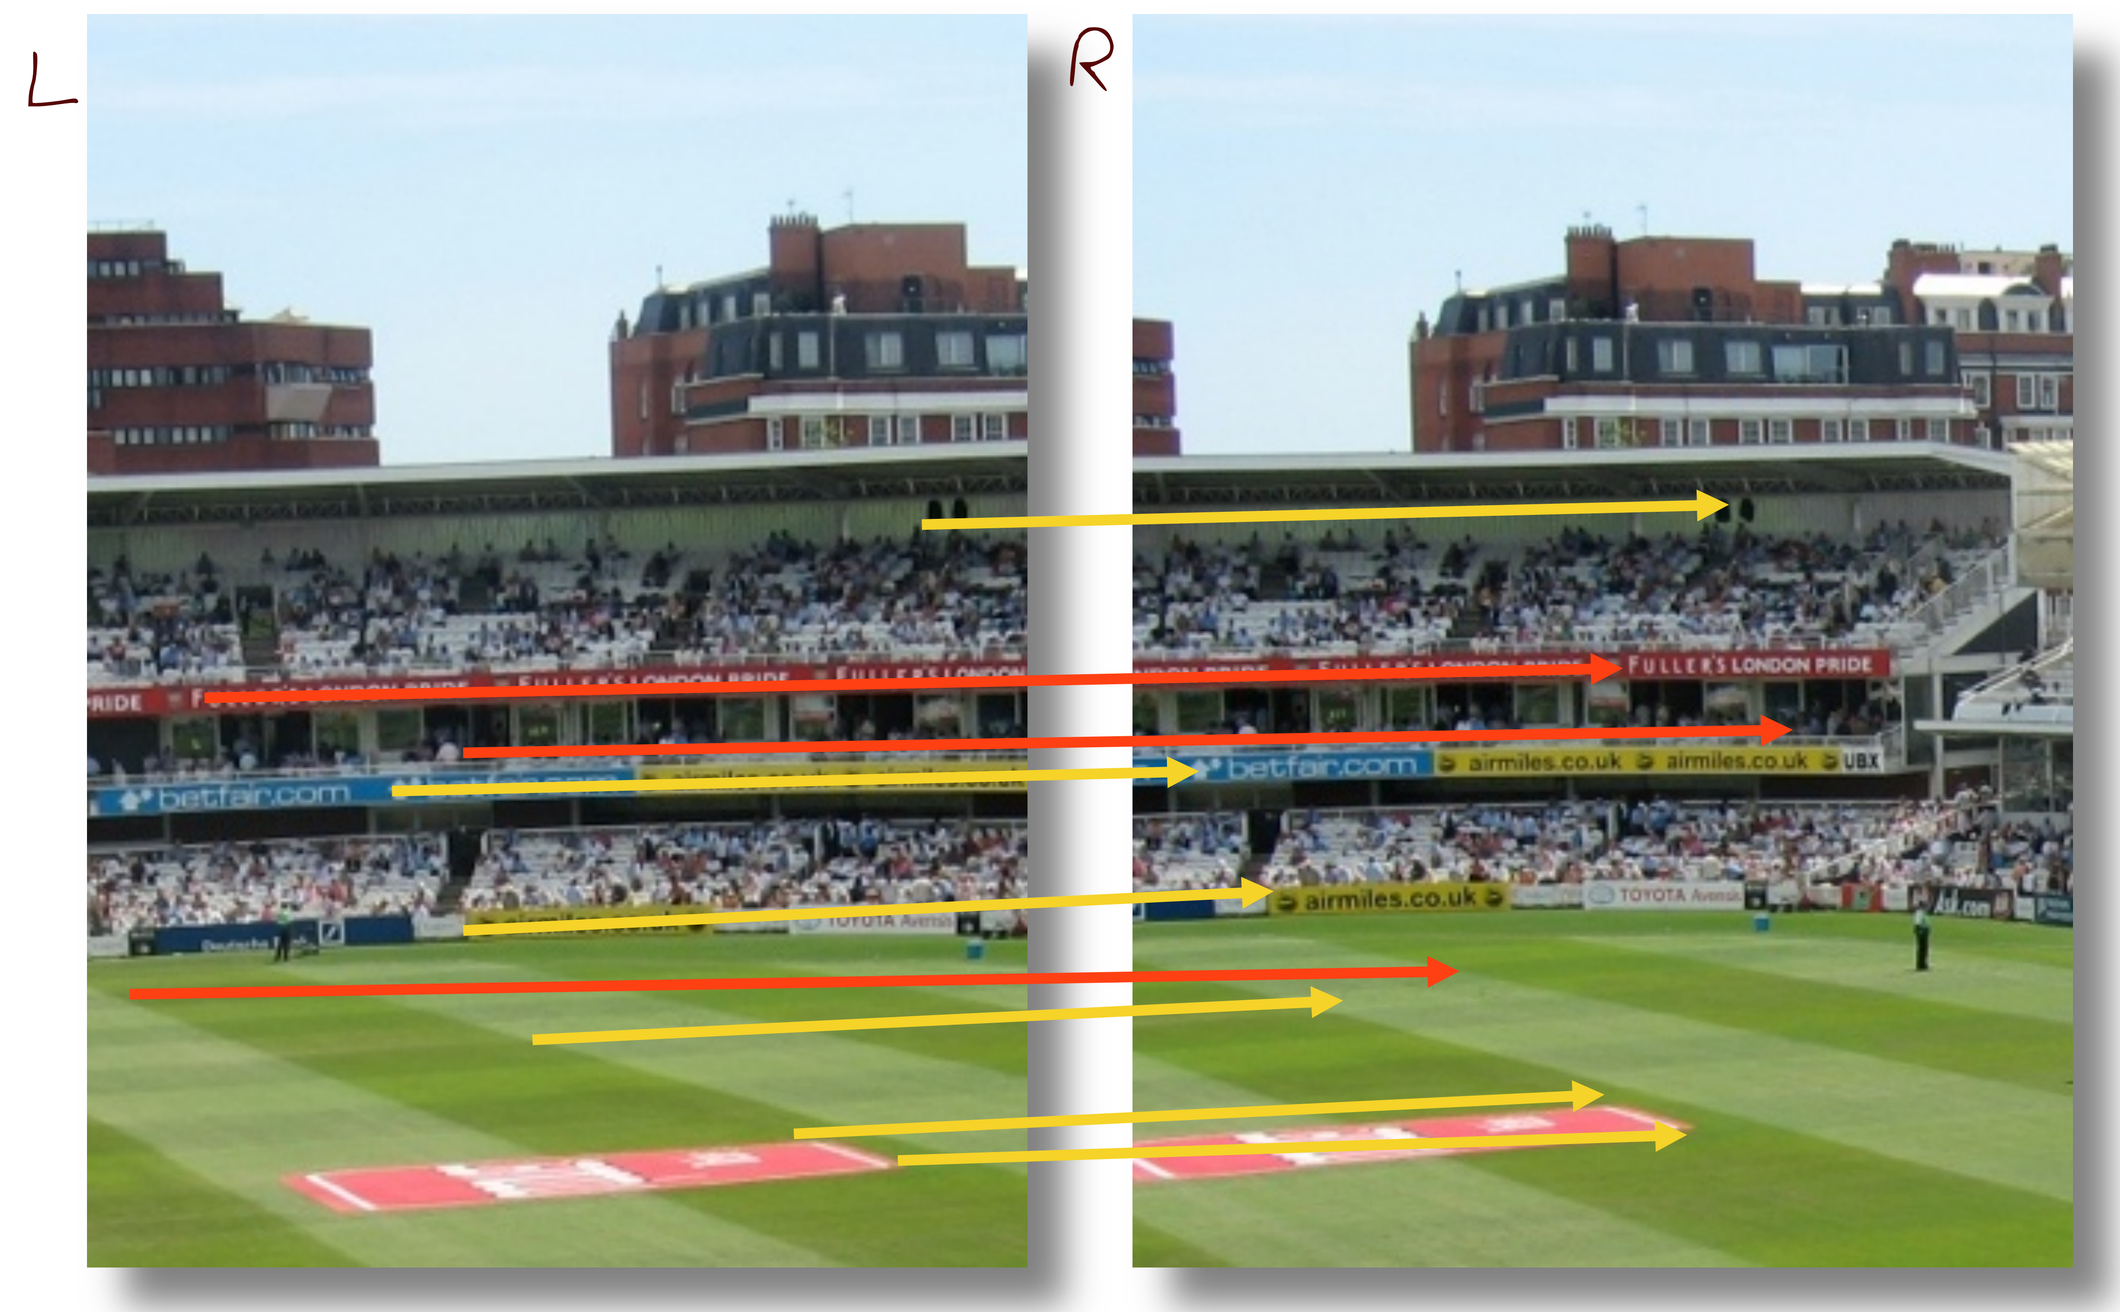

RANdom SAmple Consensus (RANSAC)

* Select ONE match, Count INLIERS
* Find "average" translation vector

RANdom SAmple Consensus (RANSAC)

Loop till find a convergence/popular H:

1. Select four feature pairs (at random)
2. Compute homography H (exact)
3. Compute inliers where: $SSD(p_{in}', H\, p_{in}) < \varepsilon$
4. Keep largest set of inliers
5. Re-compute least-squares H estimate on all of the inliers

Key idea: Not that there are more inliers than outliers, but that the outliers are wrong in different ways.

Not Just a Plane

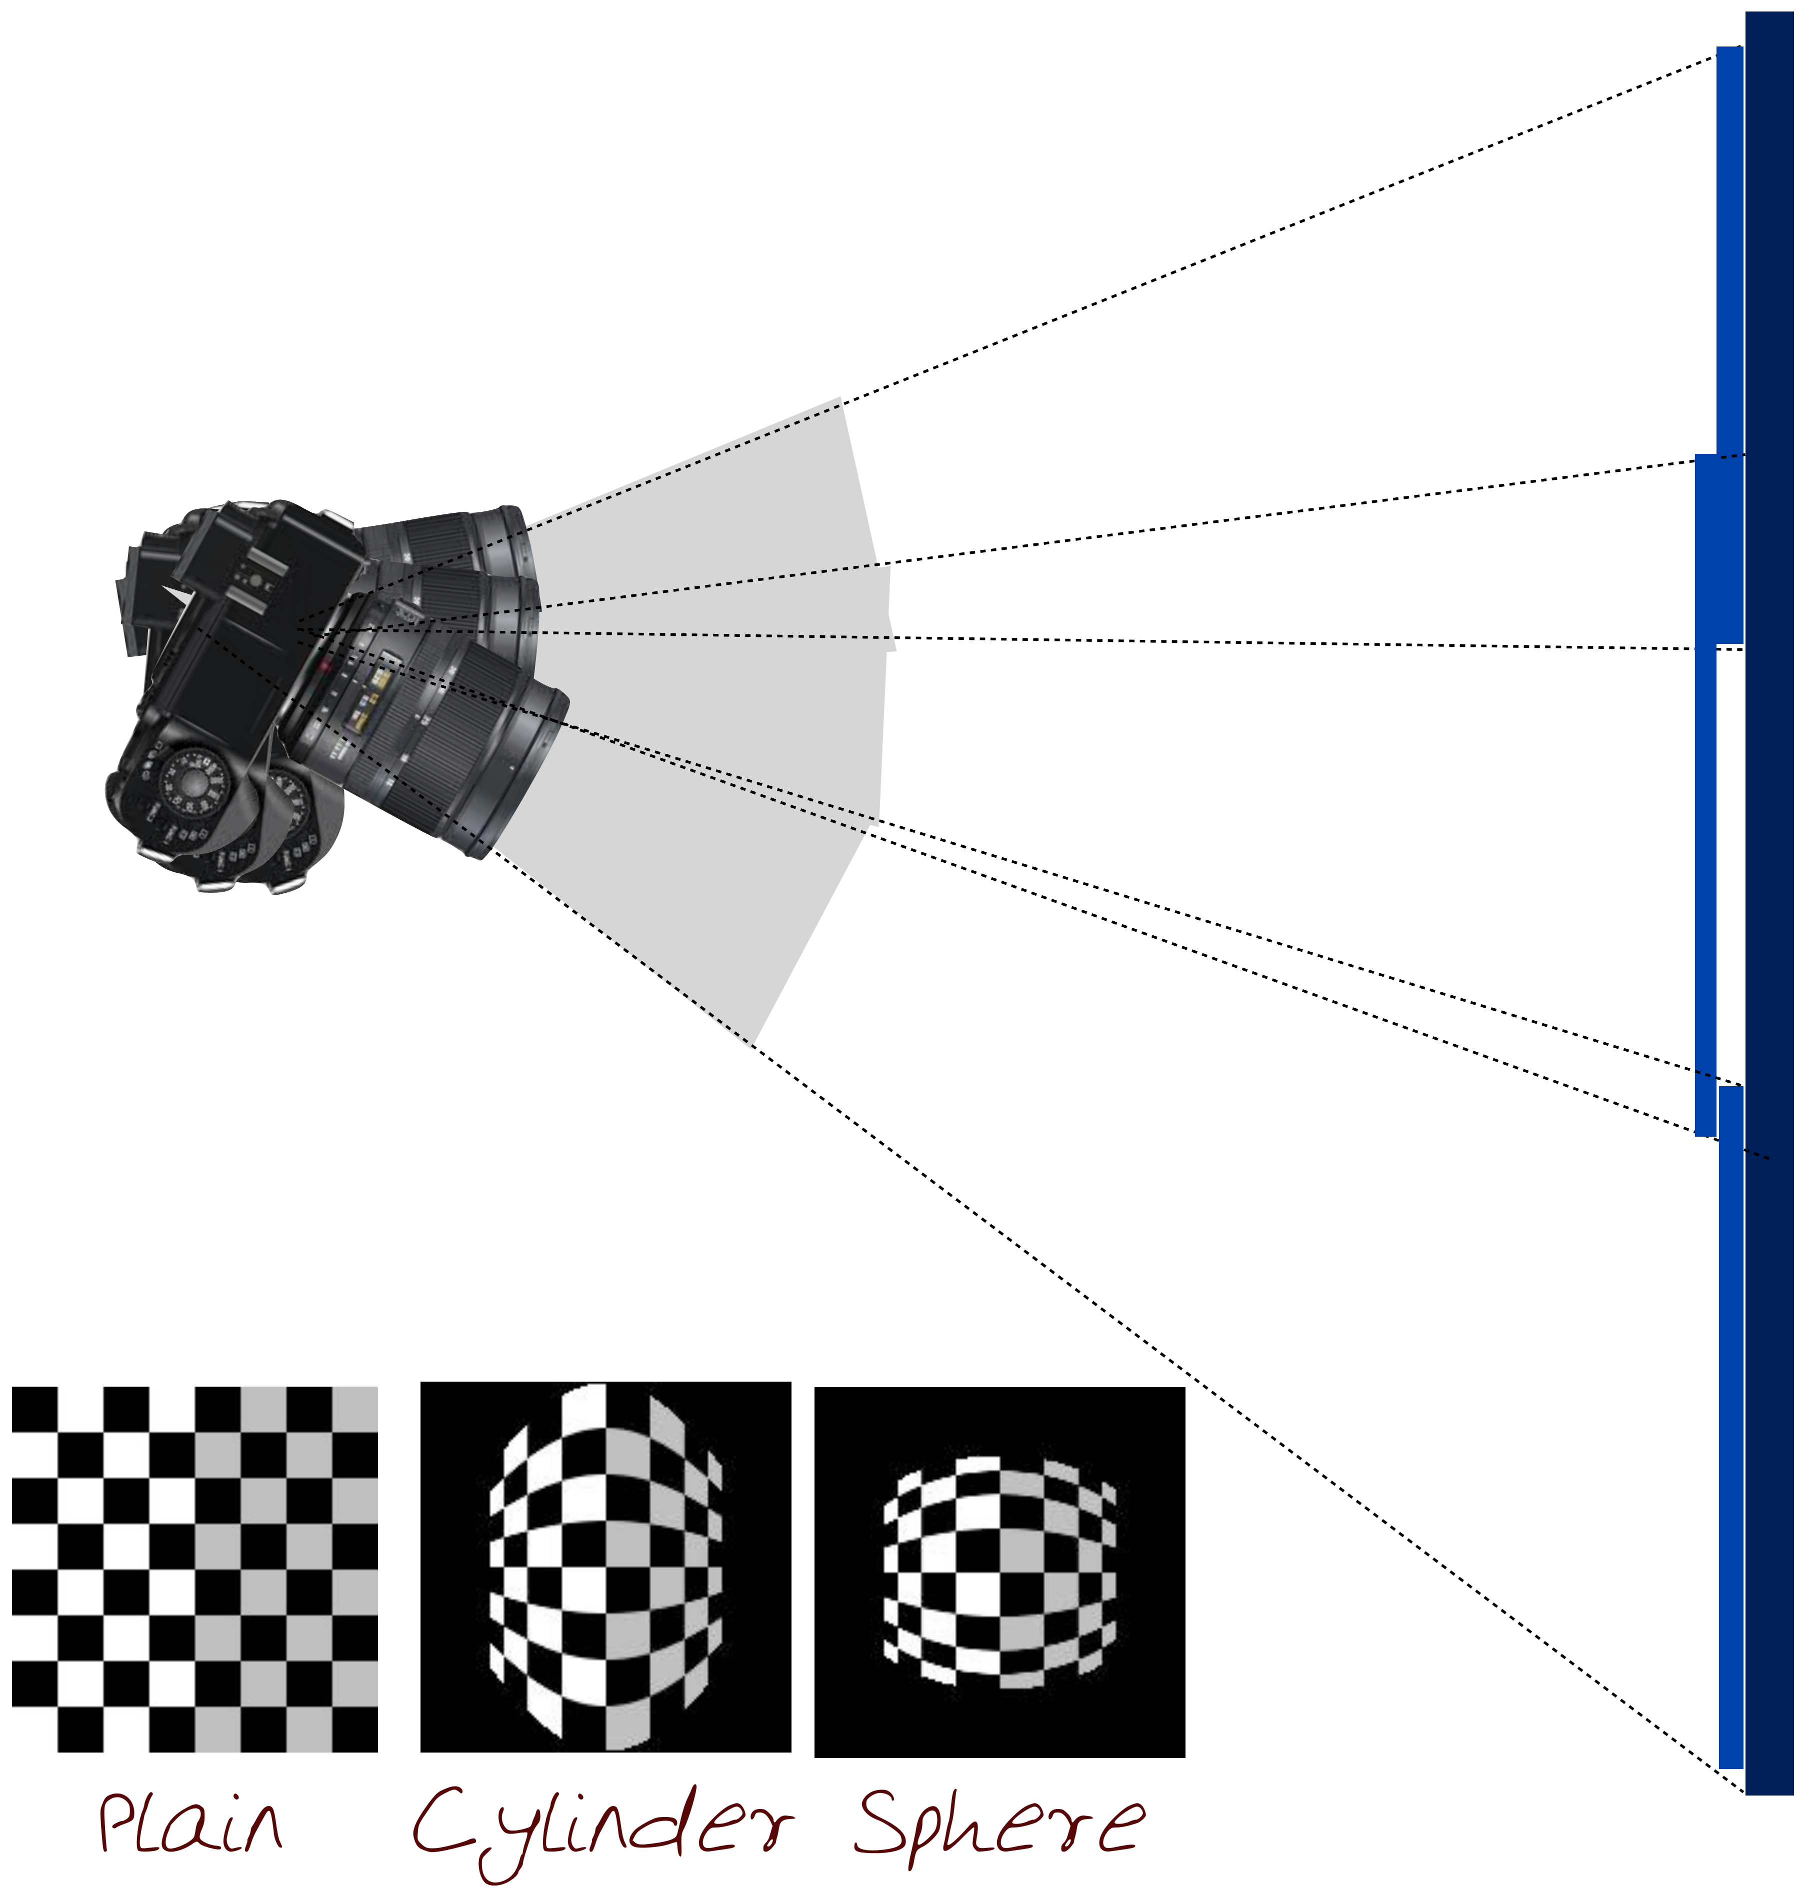

Projection Plane

* Cylinder
* Sphere

Planar, Spherical, Cylindrical Panoramas

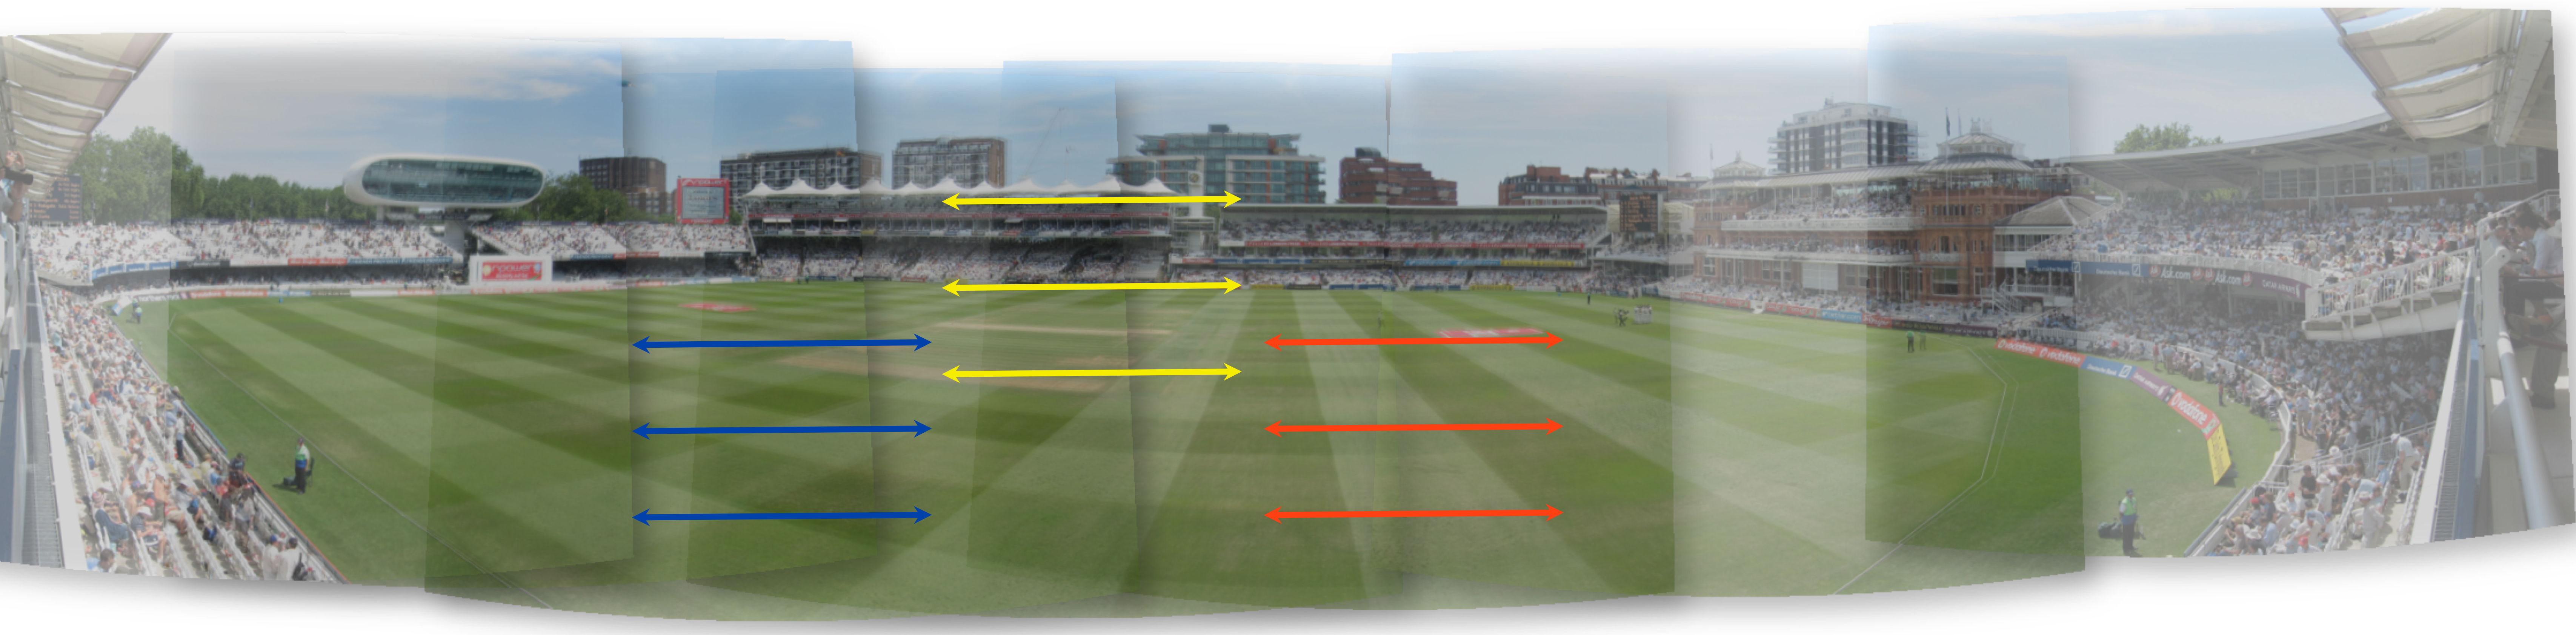

"Finding Panoramas"

Using RANSAC and related matching techniques, we can find images next to each other that form a panorama. So we don't have to take pictures in a sequence.

Brown and Lowe (2003).

Further Reading

* Brown and Lowe (2003). "Recognising Panoramas." International Conference on Computer Vision (ICCV2003) (pdf | bib | ppt)
* Microsoft Research Image Composite Editior (ICE)
* Panorama Tools Graphical User Interface (PTGui)
* Hugin Panorama Photo Stitcher

High Dynamic Range

Inside, No Lights
Long Exposure

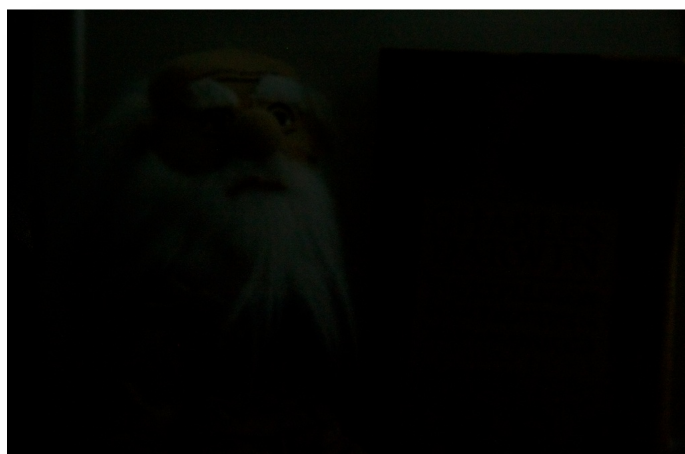

Into the Sun

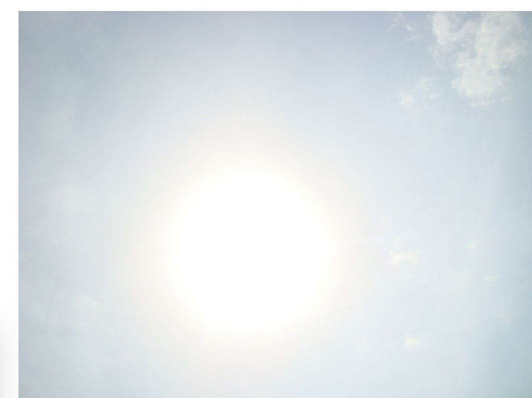

Outside,
in the Sun

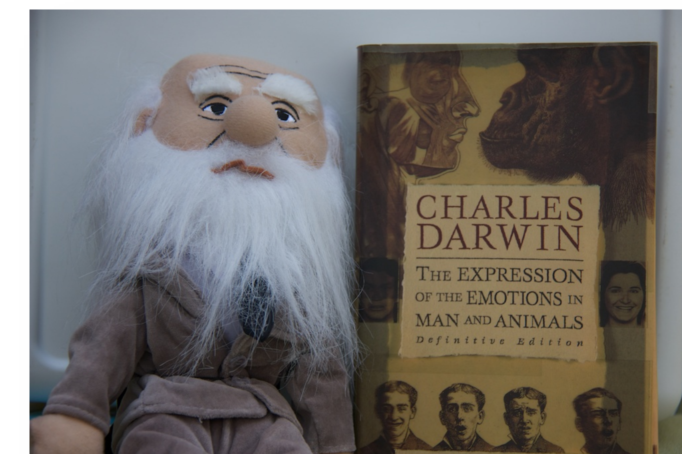

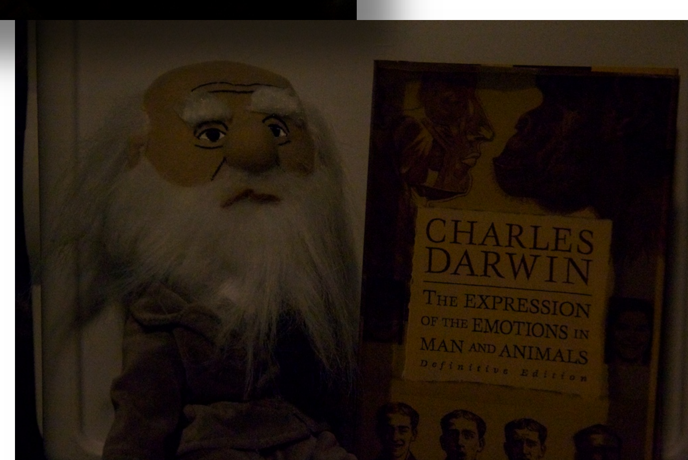

Inside,
Incandescent Light

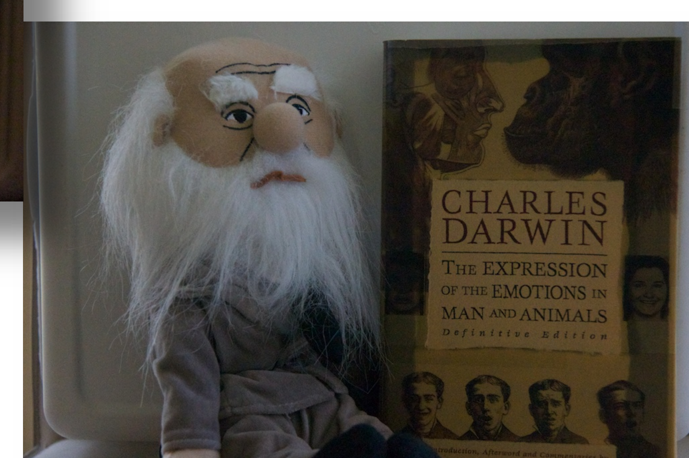

Inside, Near Window
(Natural Light)

Outside,
Under Shade

Dynamic Range

Luminance: A photometric measure of the luminous intensity per unit area of light traveling in a given direction. measured in candela per square meter (cd/m^2).

Luminance ($\log cd/m^2$)

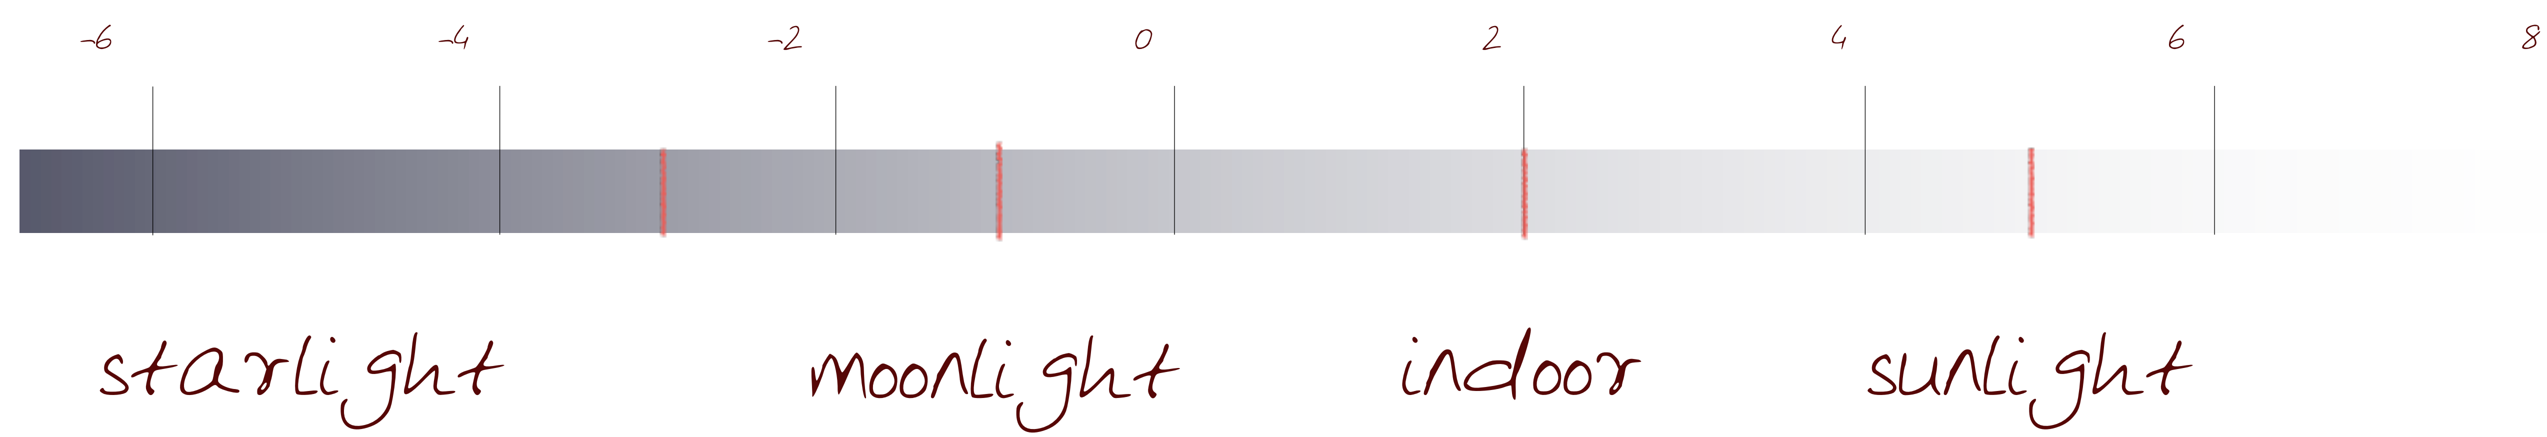

*Human Static Contrast Ratio: 100:1 (10^2:1) ➔ about 6.5 f-stops

*Human Dynamic Contrast Ratio: 1,000,000:1 (10^6:1) ➔ about 20 f-stops

Limited Dynamic Range of Current Cameras


Short Exposure: Snow and Outside Visible


Long Exposure: Inside Visible

* Need about 5-10 million values to store all brightnesses around us
* 8-bit images provide only 256 values!!

Limited Dynamic Range of Current Cameras

Short Exposure: Snow and Outside Visible

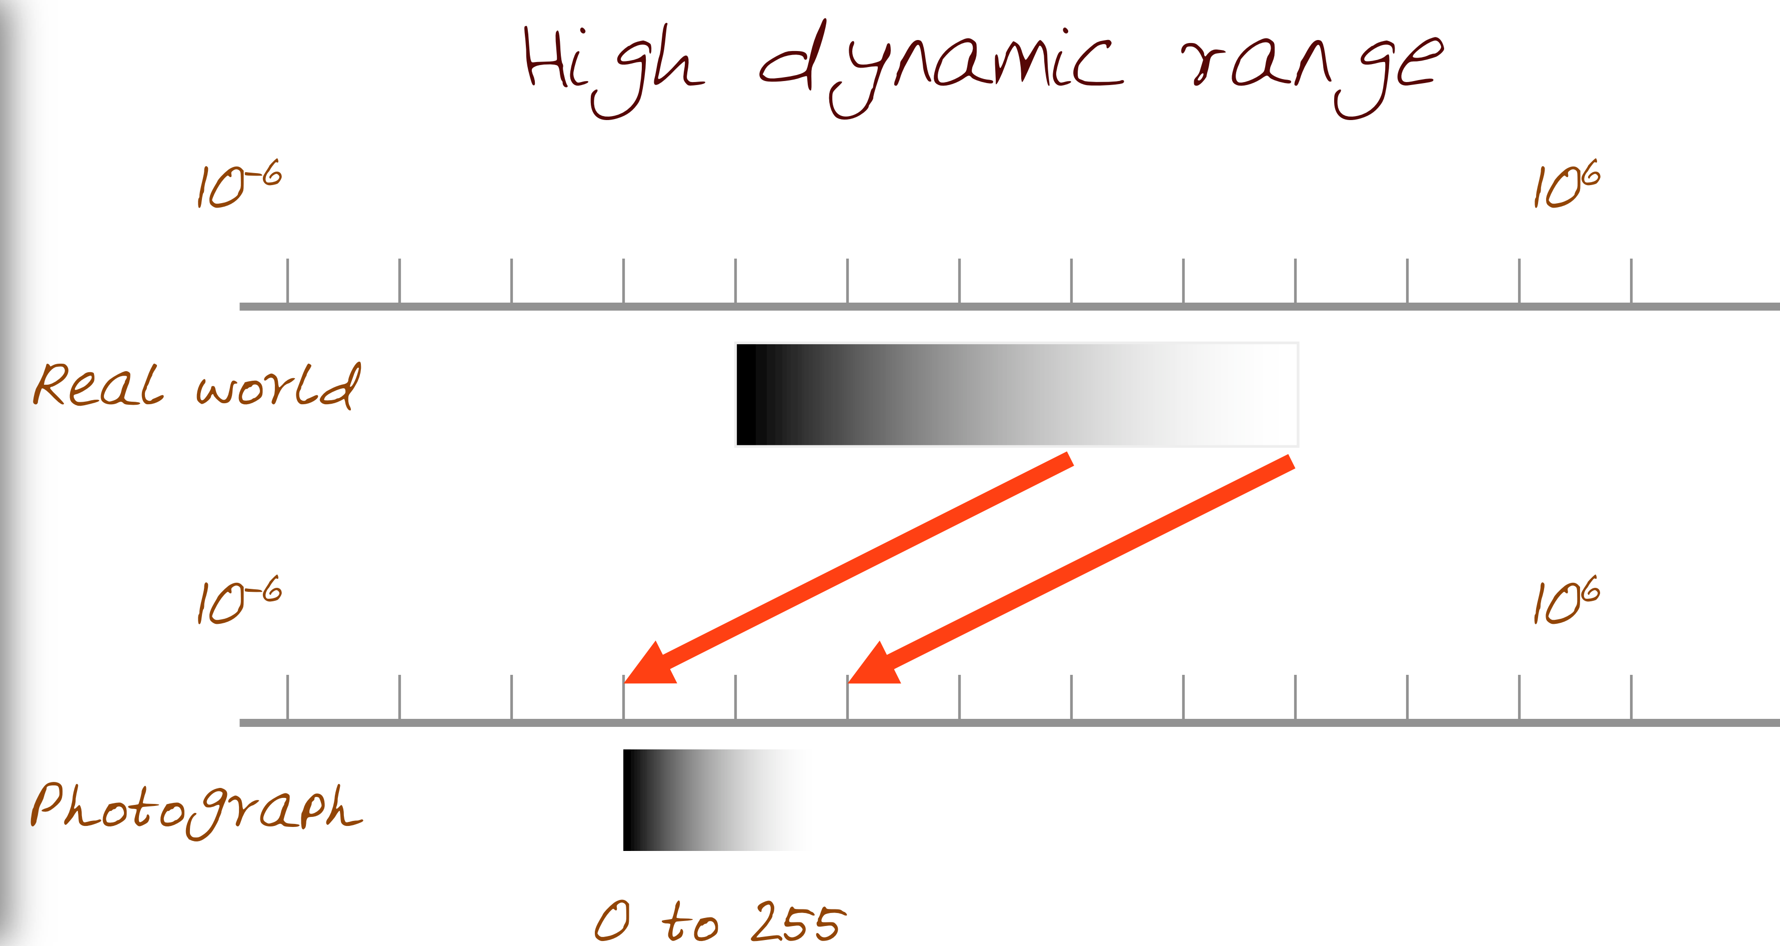

- Need about 5-10 million values to store all brightnesses around us.
- 8-bit images provide only 256 values!!

Limited Dynamic Range of Current Cameras

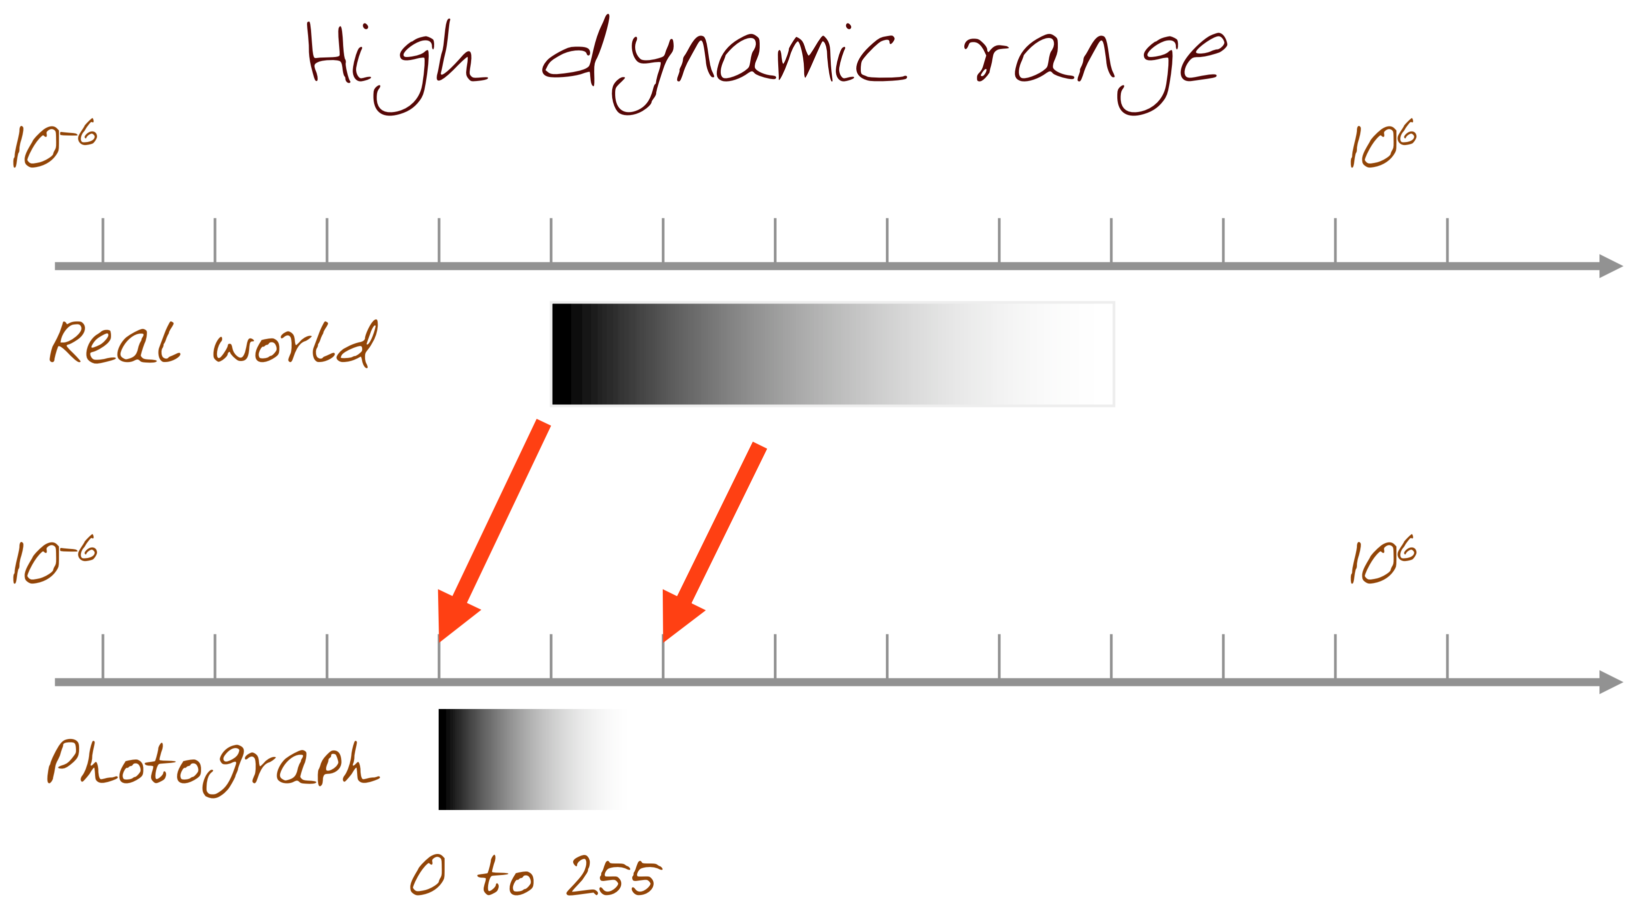

Long Exposure: Inside Visible

- Need about 5-10 million values to store all brightnesses around us.
- 8-bit images provide only 256 values!!

Relationship Between Image and Scene Brightness

The Image Acquisition Pipeline

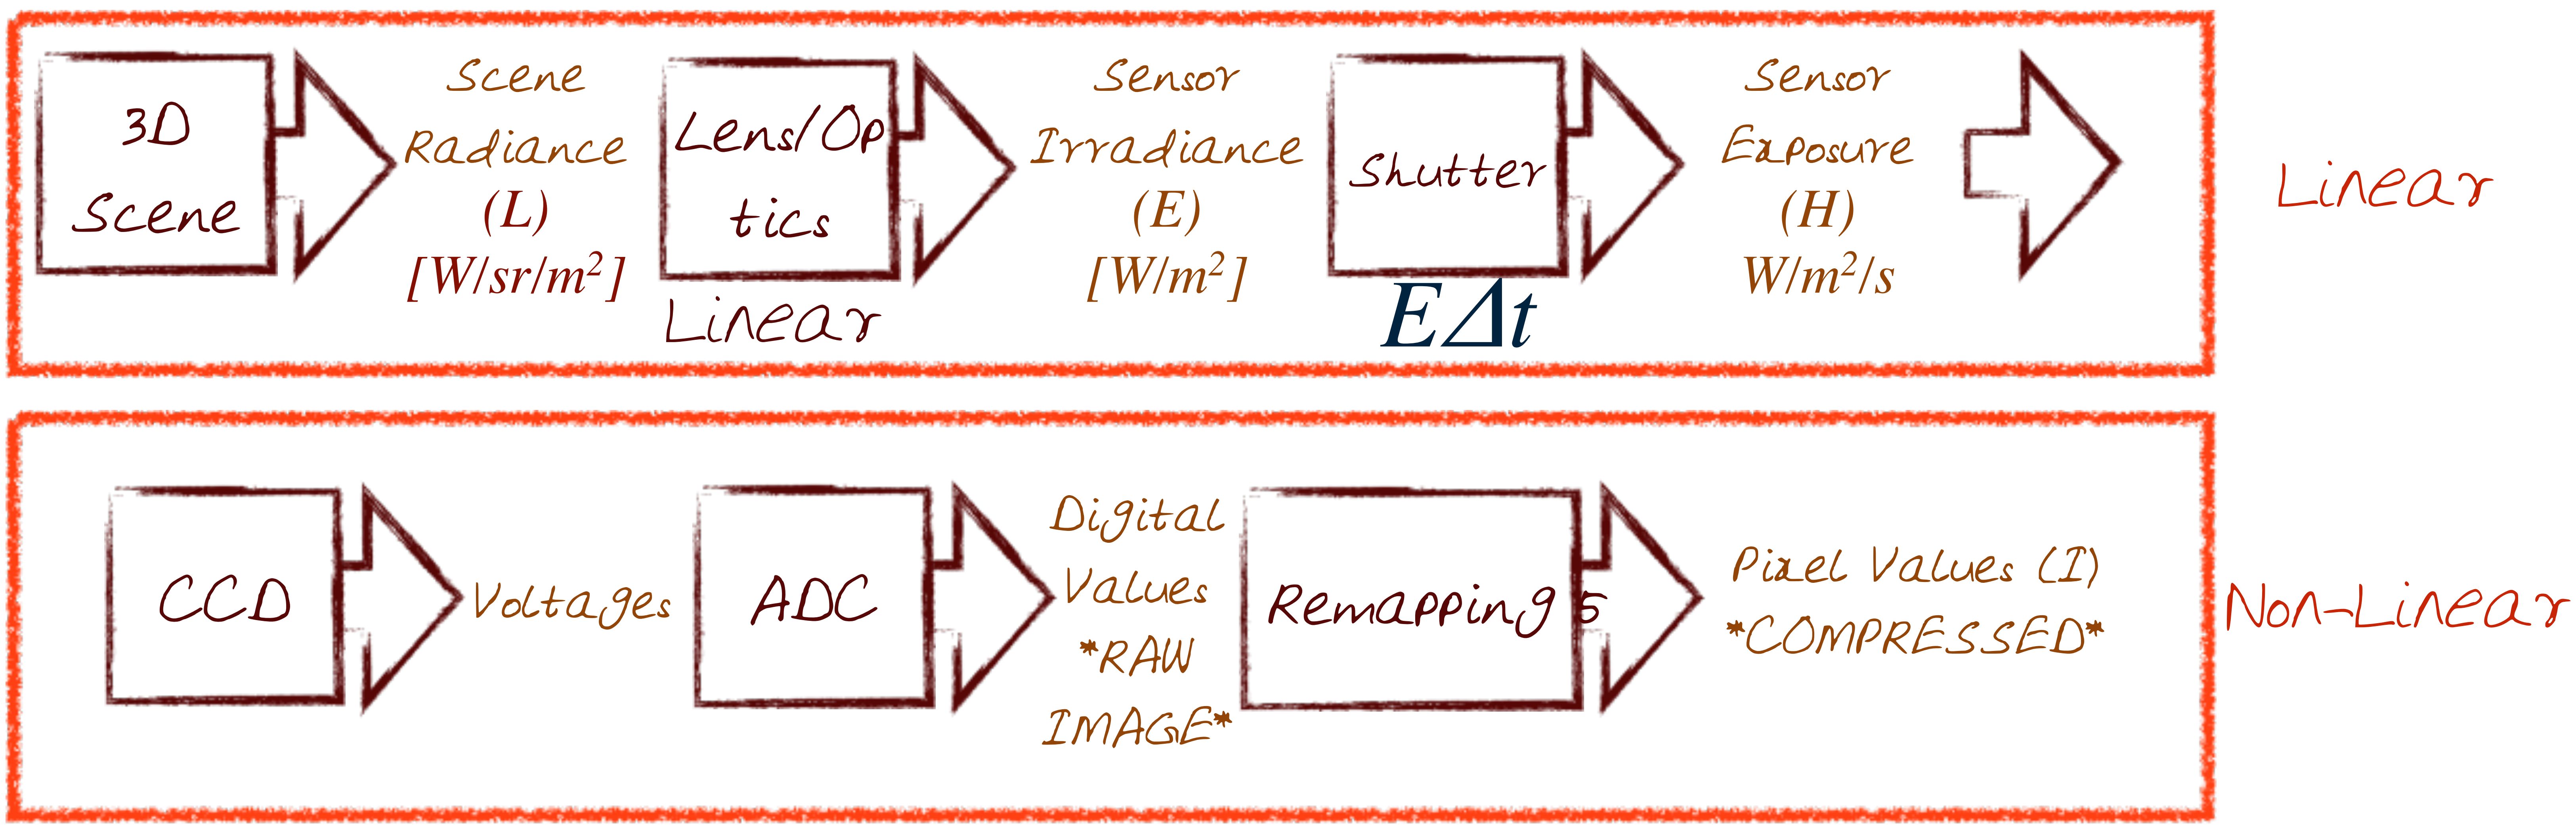

Relationship Between Image and Scene Brightness

$$g: L \rightarrow E \rightarrow H \rightarrow I \quad \Longleftrightarrow \quad g^{-1}: I \rightarrow H \rightarrow E \rightarrow L$$

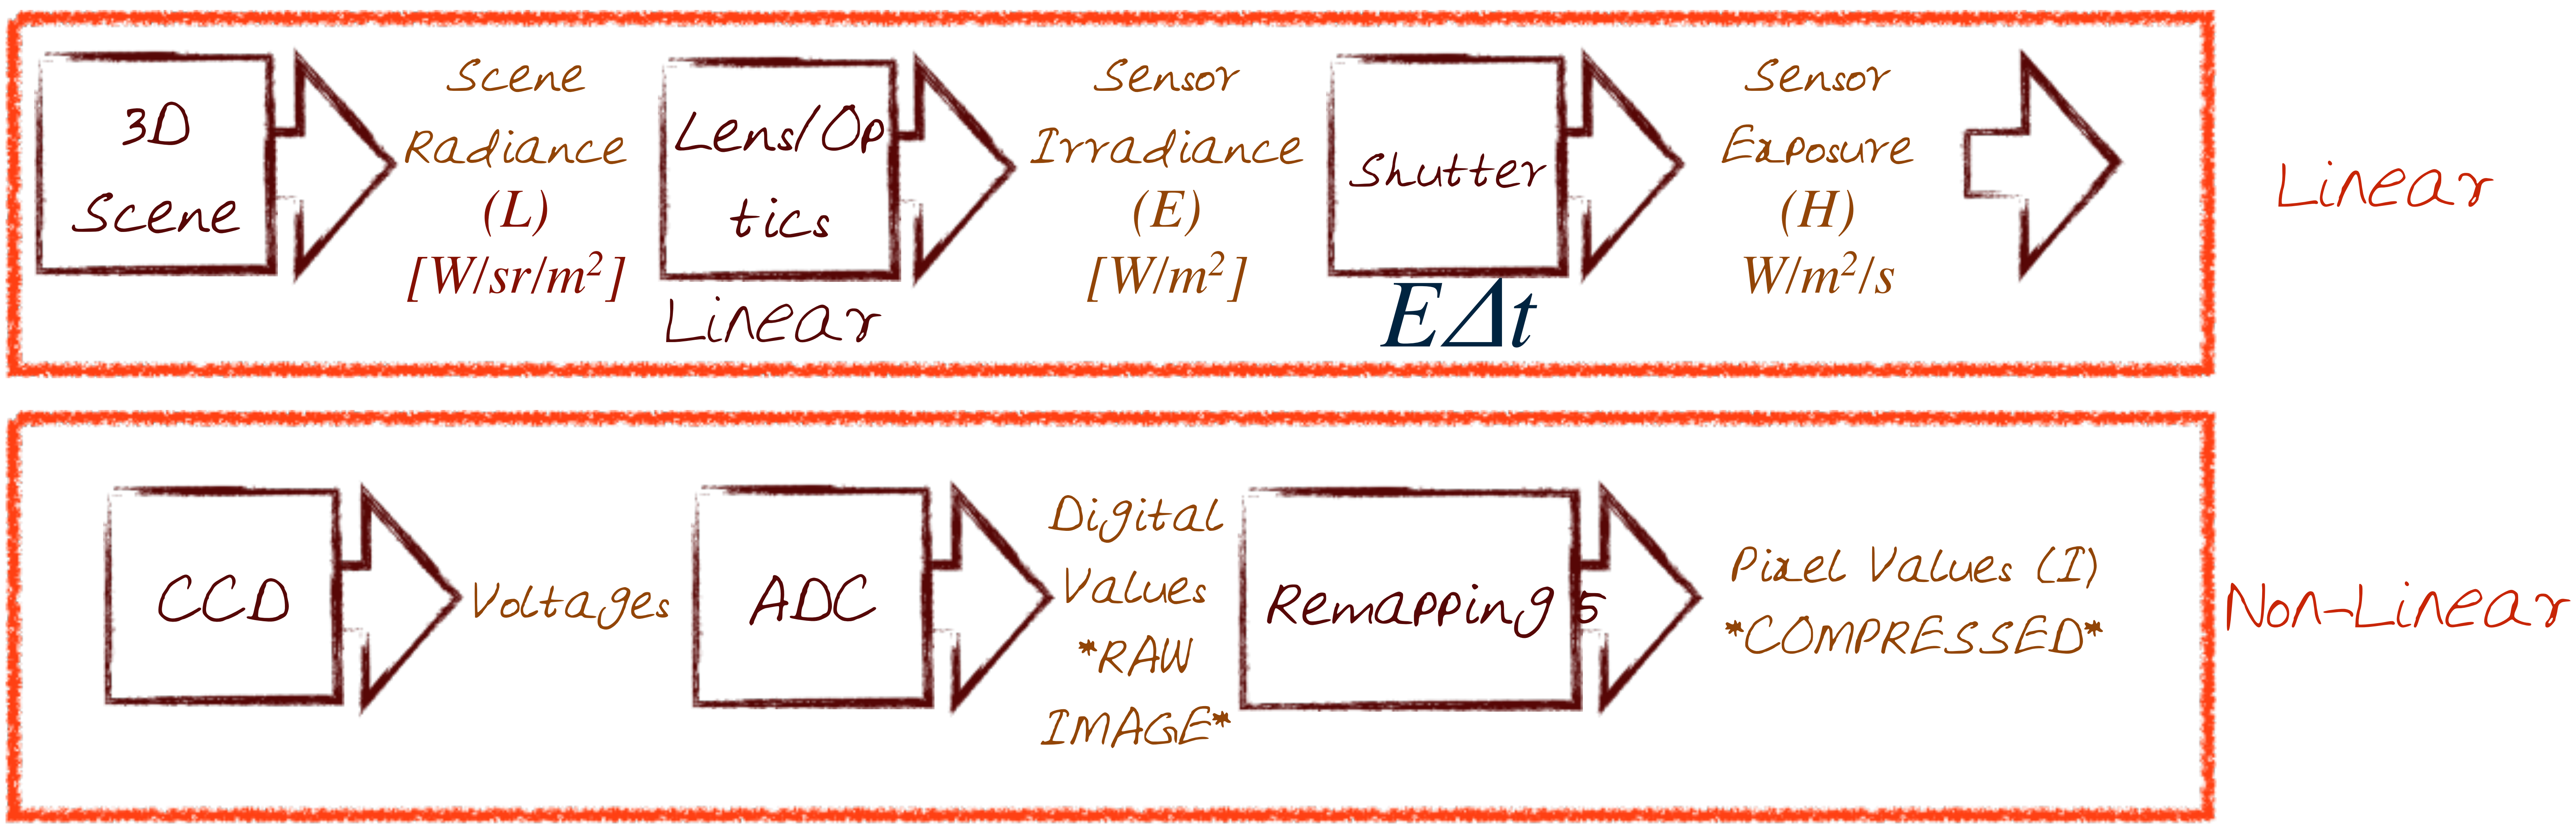

Camera Calibration

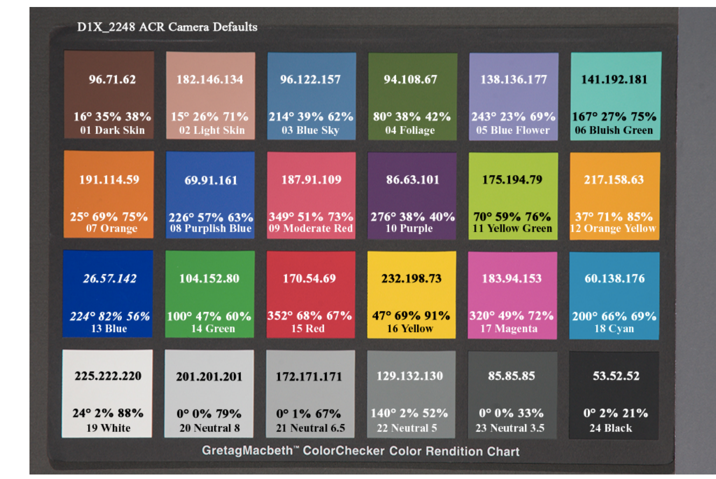

- Geometric
  - How pixel coordinates relate to directions in the world
- Radiometric / Photometric
  - How pixel values relate to radiance amounts in the world

Radiometric Calibration

$$g: L \rightarrow E \rightarrow H \rightarrow I \Leftrightarrow g^{-1}: I \rightarrow H \rightarrow E \rightarrow L$$

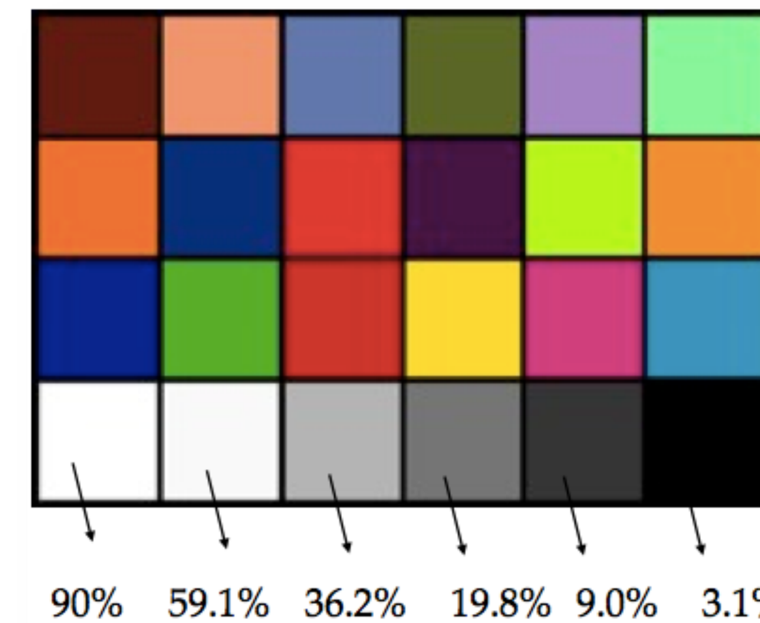

- A Color Chart with know reflectances
- Multiple camera exposures to fill up the curve
- Method assumes constant lighting on all patches and works best when source is far away (example sunlight)
- Unique inverse exists because g is monotonic and smooth for all cameras

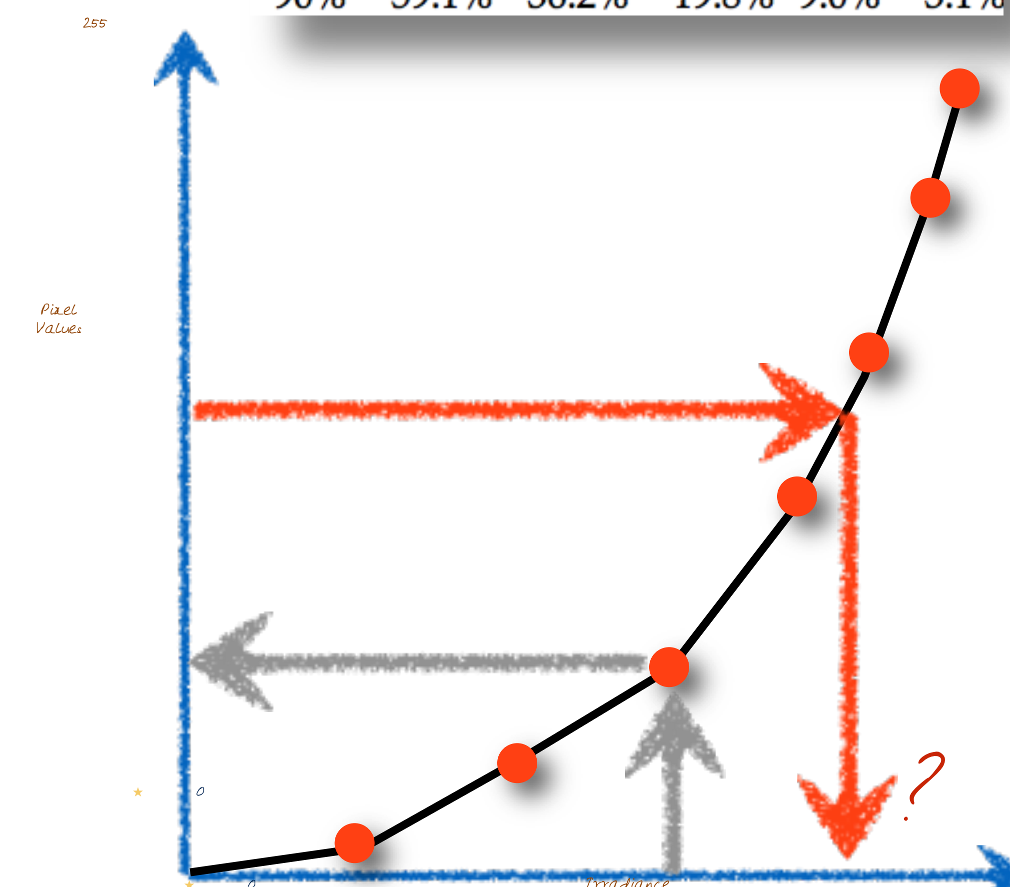

(Grossberg and Nayar 2003)

A Sequence of Images of Different Exposures

Series of Images

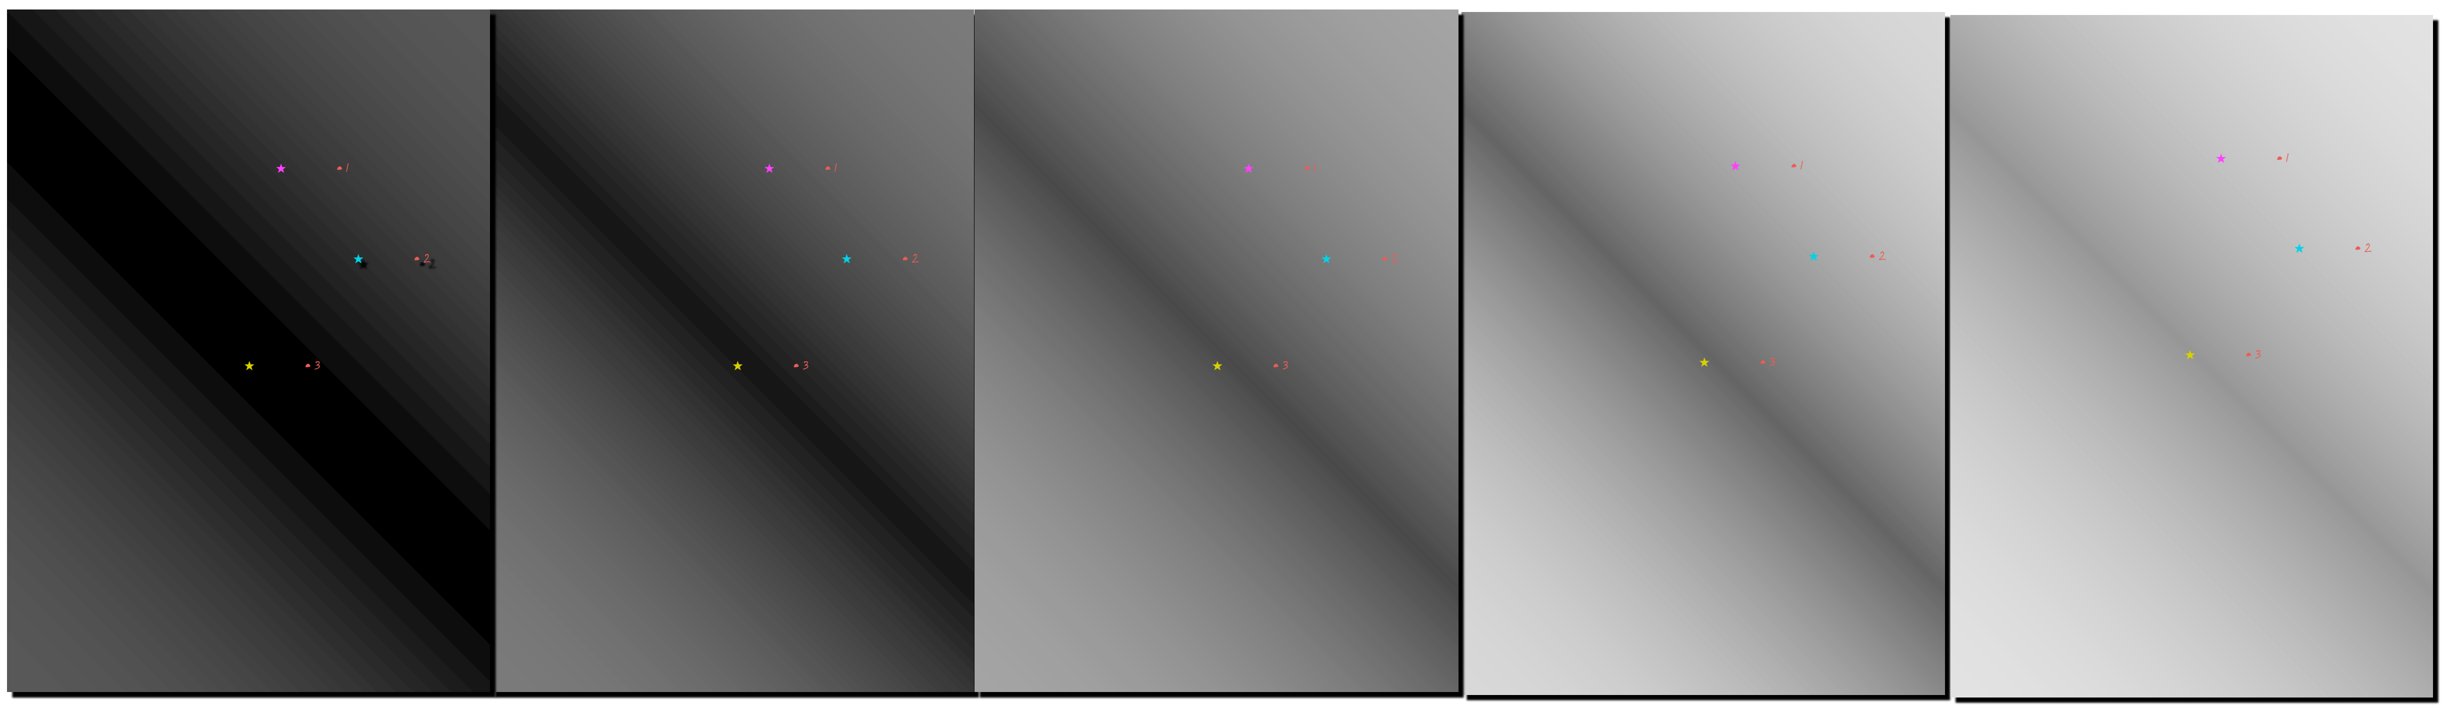

Δt = 1/64 sec   Δt = 1/16 sec   Δt = 1/4 sec   Δt = 1 sec   Δt = 4 sec

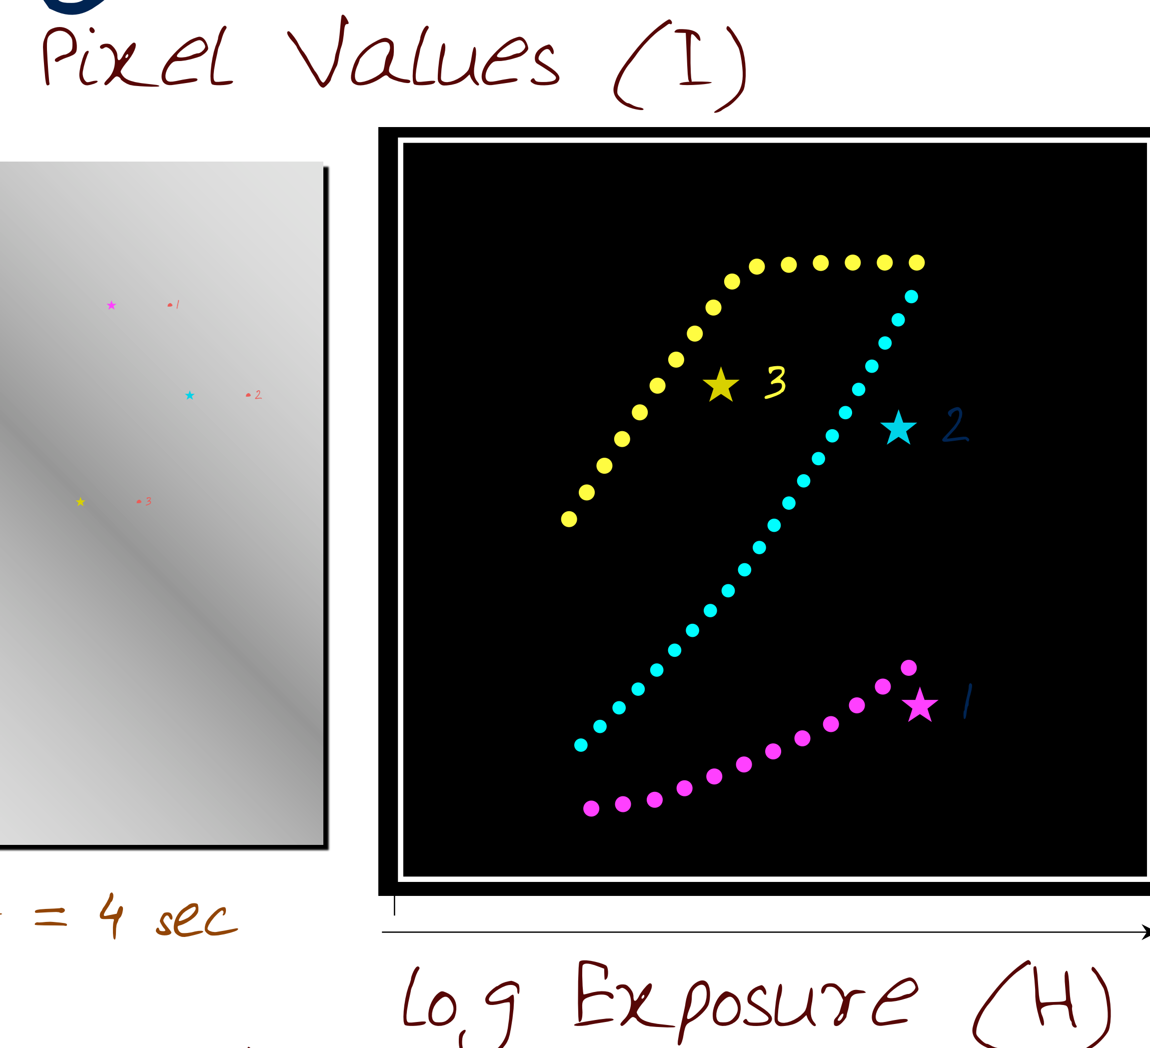

Pixel Values (I) = g(Exposure)

Exposure (H) = Irradiance (E) * Δt

log Exposure (H) = log Irradiance (E) + log Δt

Debevec and malik 1997

Response Curves

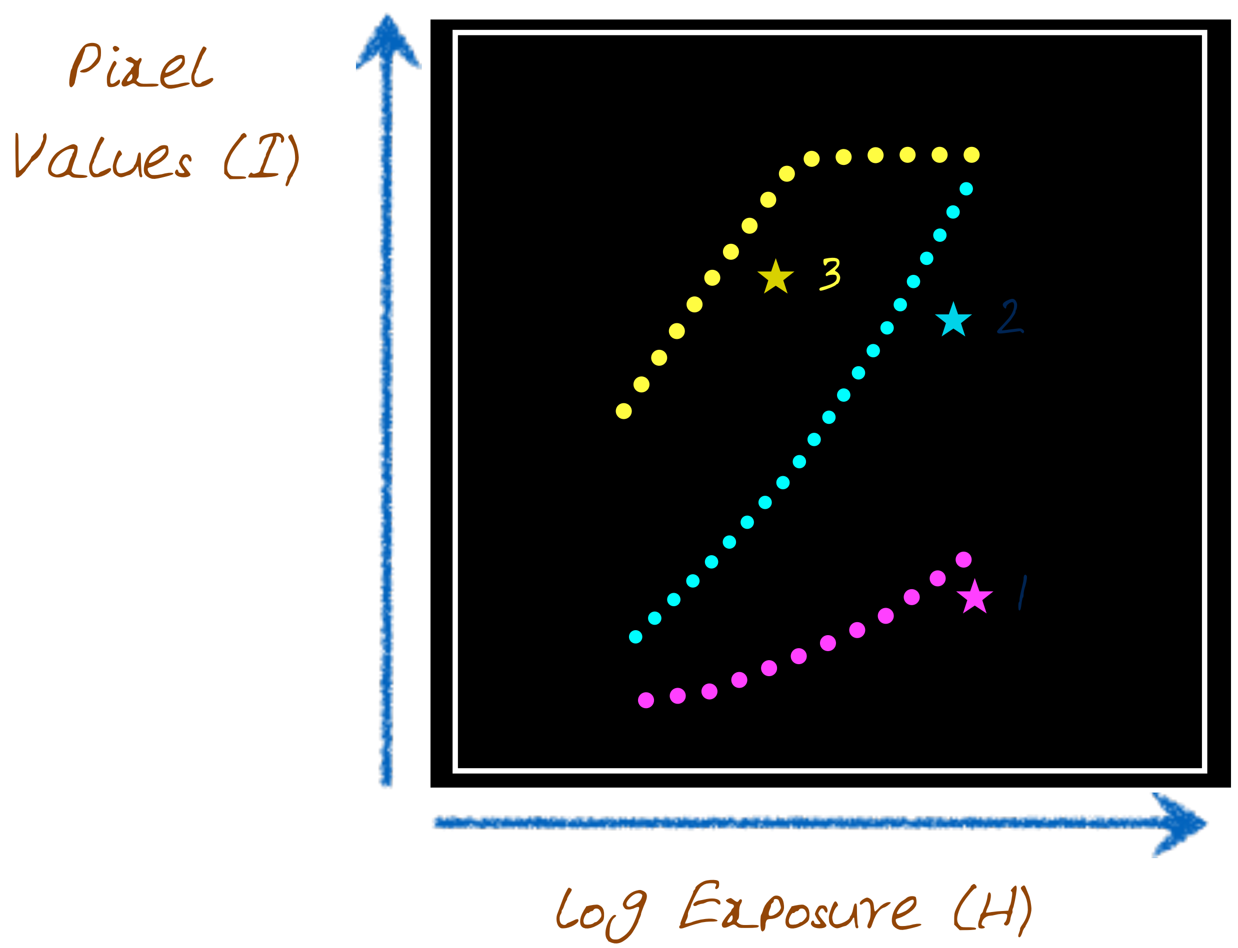

Assuming unit radiance for each pixel

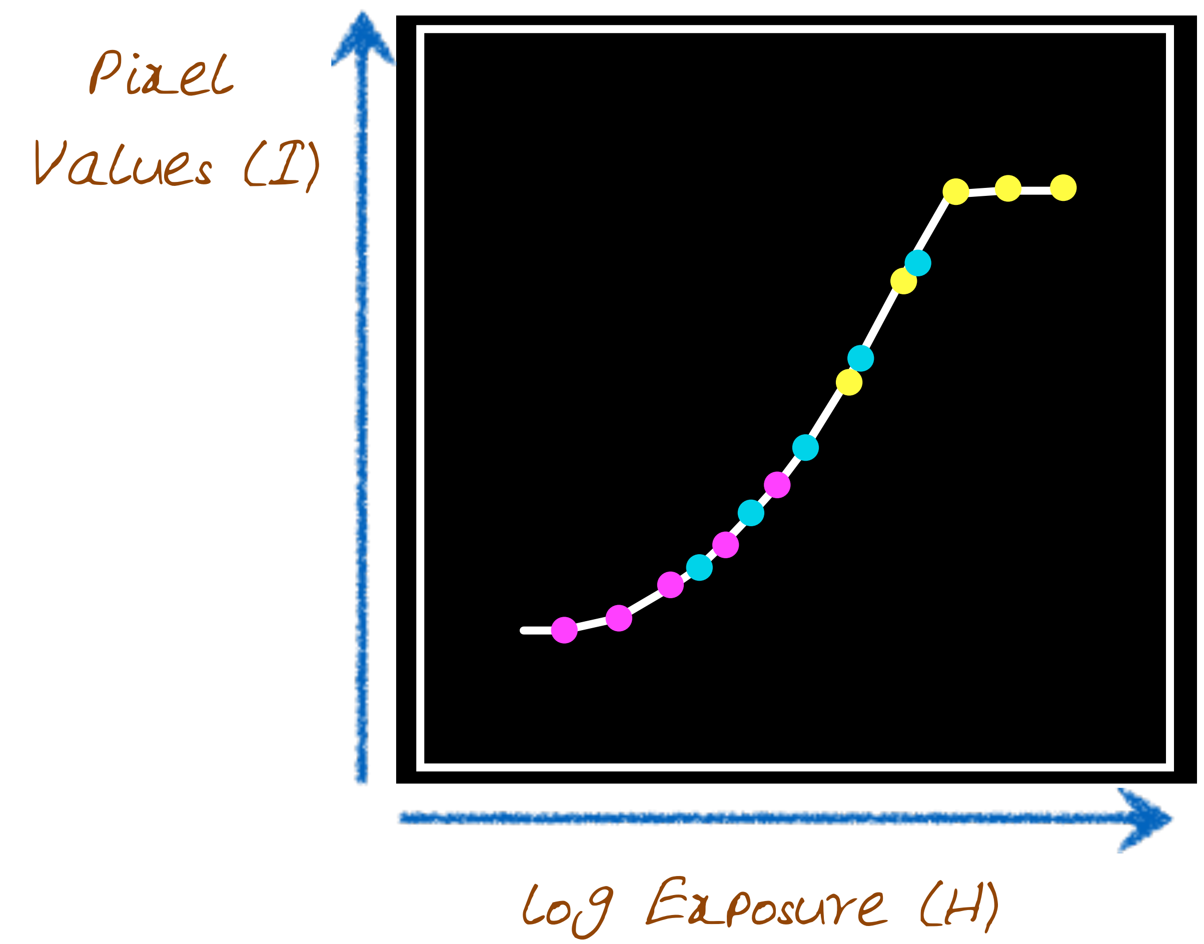

After adjusting radiances to obtain a smooth response curve

Iterative Method

* For each pixel site i in each image j, we have

$$Z_{ij} = f(X_{ij}) = f(I_i \Delta t_j) \qquad (1)$$

* So, if have f, we can estimate the irradiance image I as

$$I_i = \frac{1}{m} \sum_{j=0}^{m-1} \frac{f^{-1}(Z_{ij})}{\Delta t_j} \qquad (2)$$

* Note f^{-1} is a lookup table
* We can re-estimate as: $f^{-1}(Z_{ij}) = I_i \Delta t_j = X_{ij}$ (3)
* Iterate (2) and (3), until convergence.

Demo-time!

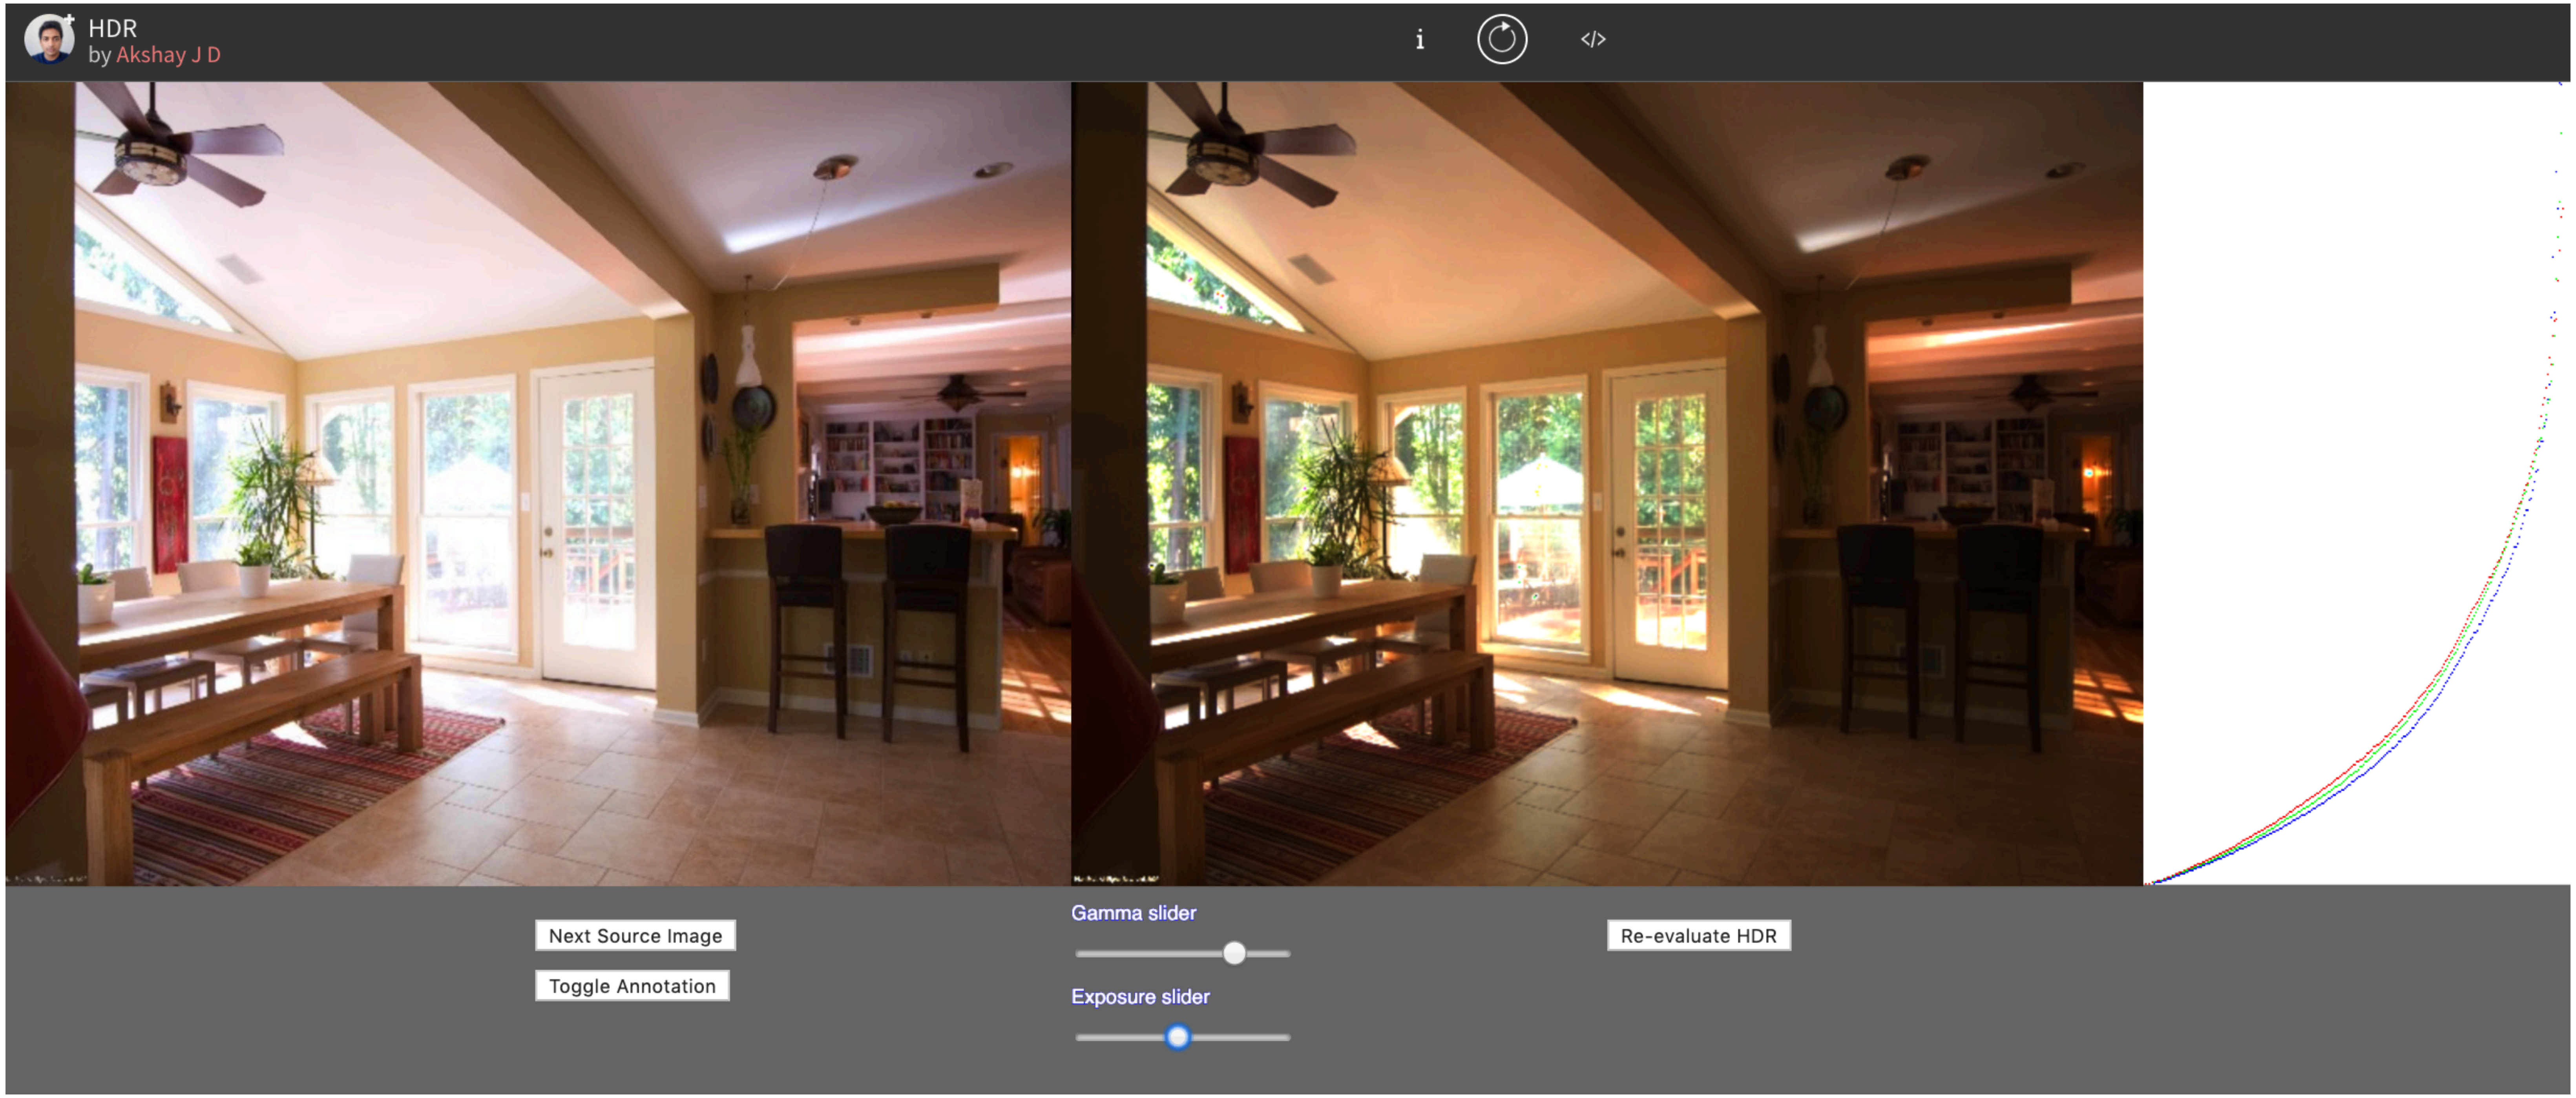
HDR
by Akshay J D
Next Source Image
Toggle Annotation
Gamma slider
Exposure slider
Re-evaluate HDR

How to Compute: Debevec 99

- Let g(z) be the discrete inverse response function
- For each pixel site i in each image j, compute

$$\ln(E_i) + \ln(\Delta t_j) = g(Z_{ij})$$

- Solve the overdetermined linear system for N pixels over P different exposure images:

Fitting Term

Smoothness Term

$$\sum_{i=1}^{N}\sum_{i=1}^{P}\left[\ln(E_i) + \ln(\Delta t_j) - g(Z_{ij})\right]^2 + \lambda \sum_{z=Z_{min}}^{Z_{max}} g''(z)^2$$

See Debevec and Malik (1997) for more details

Response Curves

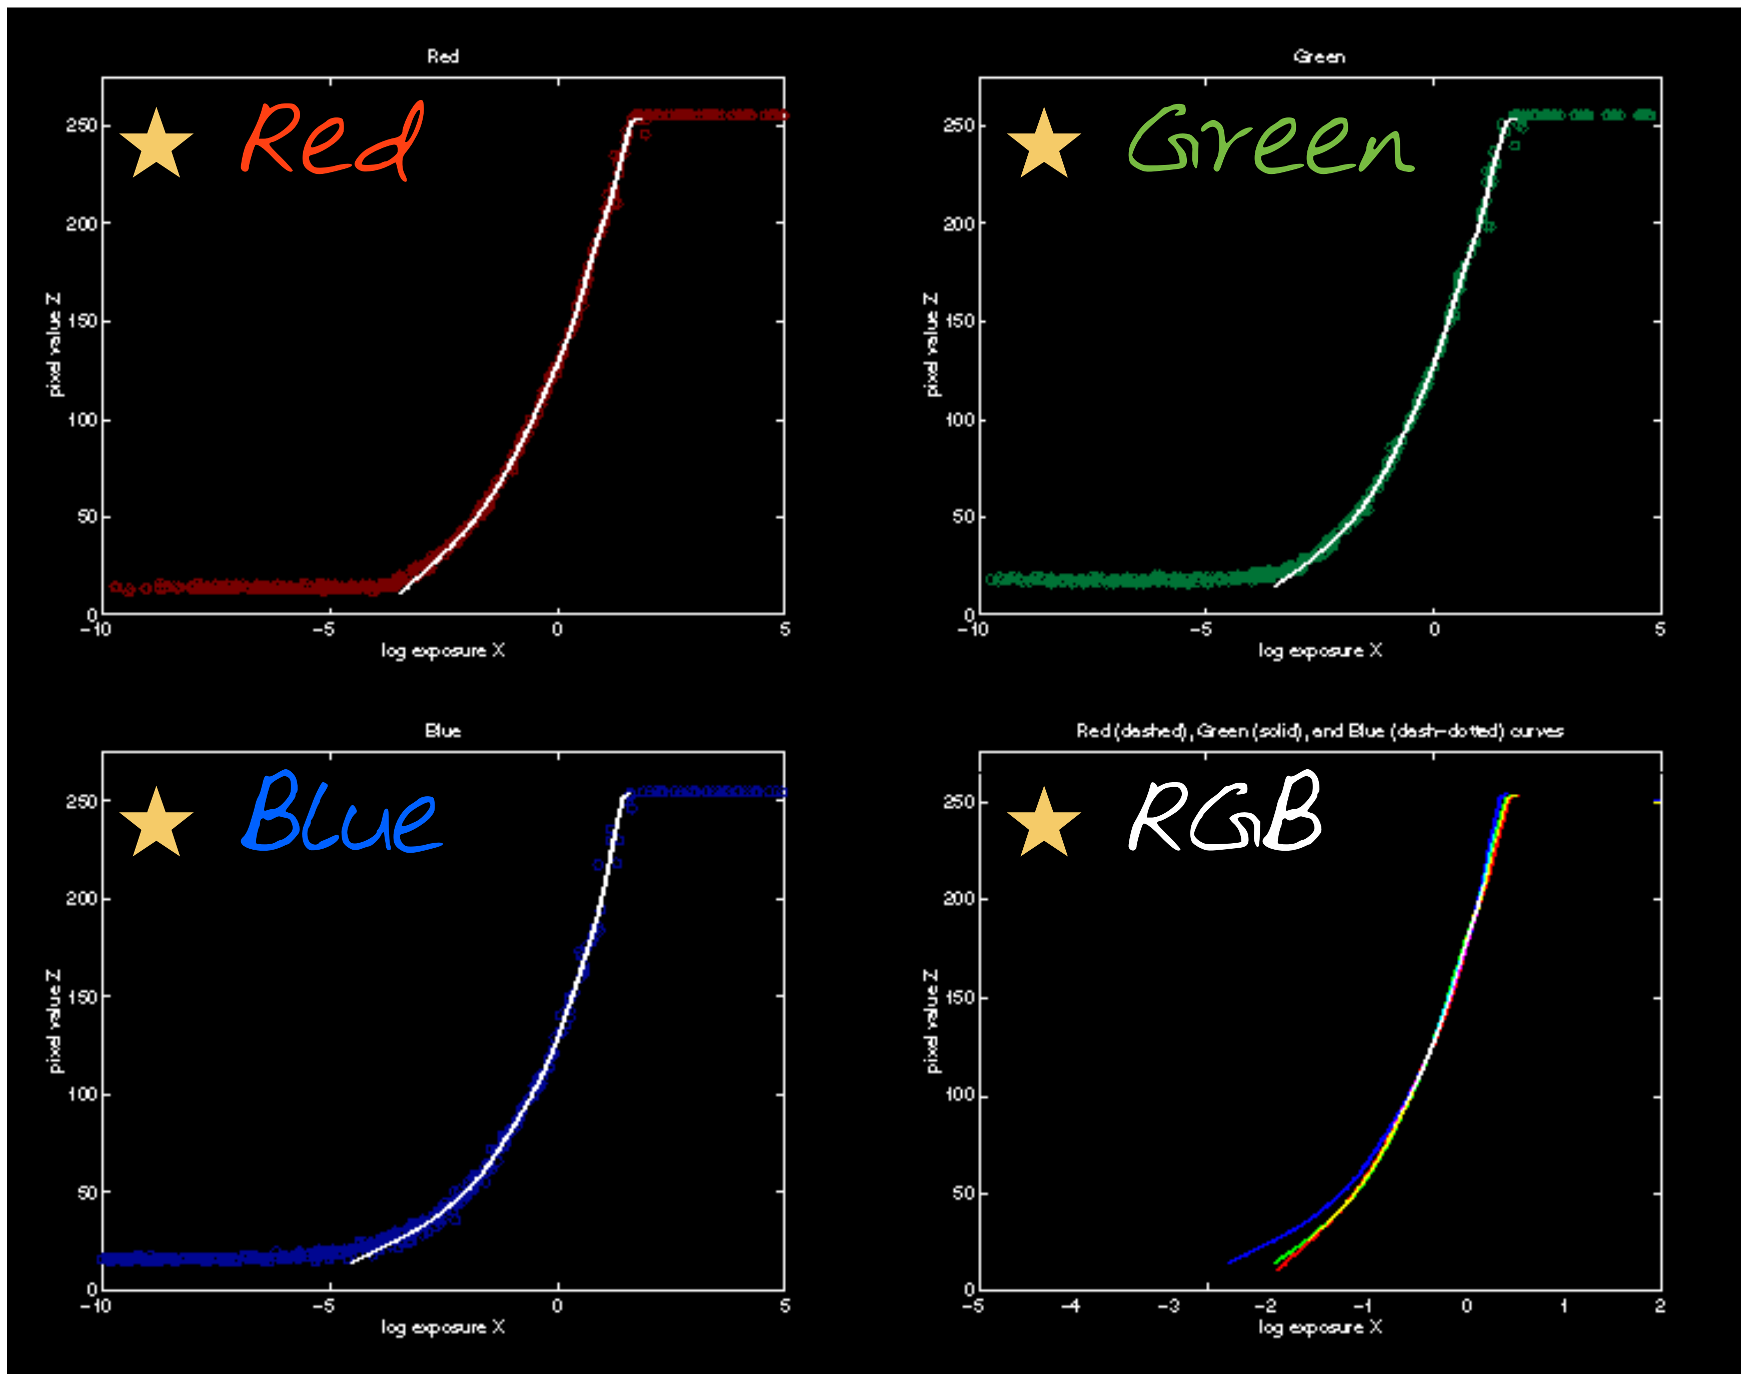

(Not actual curves for these images, used here just for demonstration)

Radiance Map

Need a New File Format

Radiance Format

32 bits / pixel

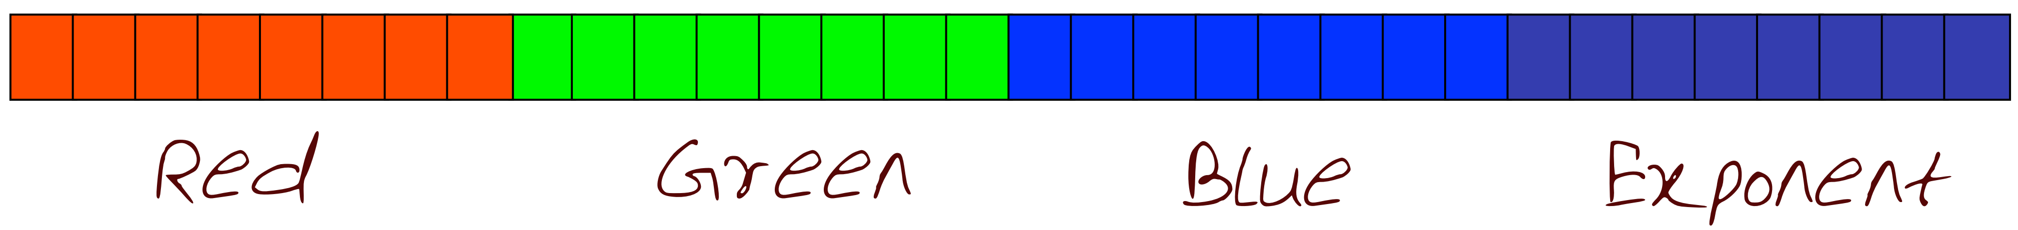

- (145, 215, 87, 149) =
- (145, 215, 87) * 2^(149-128) =
- (1190000, 1760000, 713000)

- (145, 215, 87, 103) =
- (145, 215, 87) * 2^(103-128) =
- (0.00000432, 0.00000641, 0.00000259)

Ward (2001), There are many other formats too

Now to Display it!

Tone Mapping

* map one set of colors to another
* Displaying on a medium that has limited dynamic range
  * Printers, monitors, and projectors all have a limited dynamic range
  * Inadequate to reproduce the full range of light intensities present in natural scenes

★ http://commons.wikimedia.org/wiki/File:Kanitz-Kyawsche_Gruft_in_Hainewalde_HDR.jpg

Tone Mapping

- Addresses the problem of
  - strong contrast reduction from the scene radiance to the displayable range
  - preservs the image details and color appearance
- many well-known Algorithms exist for this
- See Banterle, et al. (2011), Reinhard et al. (2002) and Durand and Dorsey (2002)

http://commonswikimedia.org/wiki/File:Kanitz-Kyawsche_Gruft_in_Hainewalde_HDR.jpg

Tone Mapping

http://en.wikipedia.org/wiki/File:Dundus_Square.jpg

High dynamic range

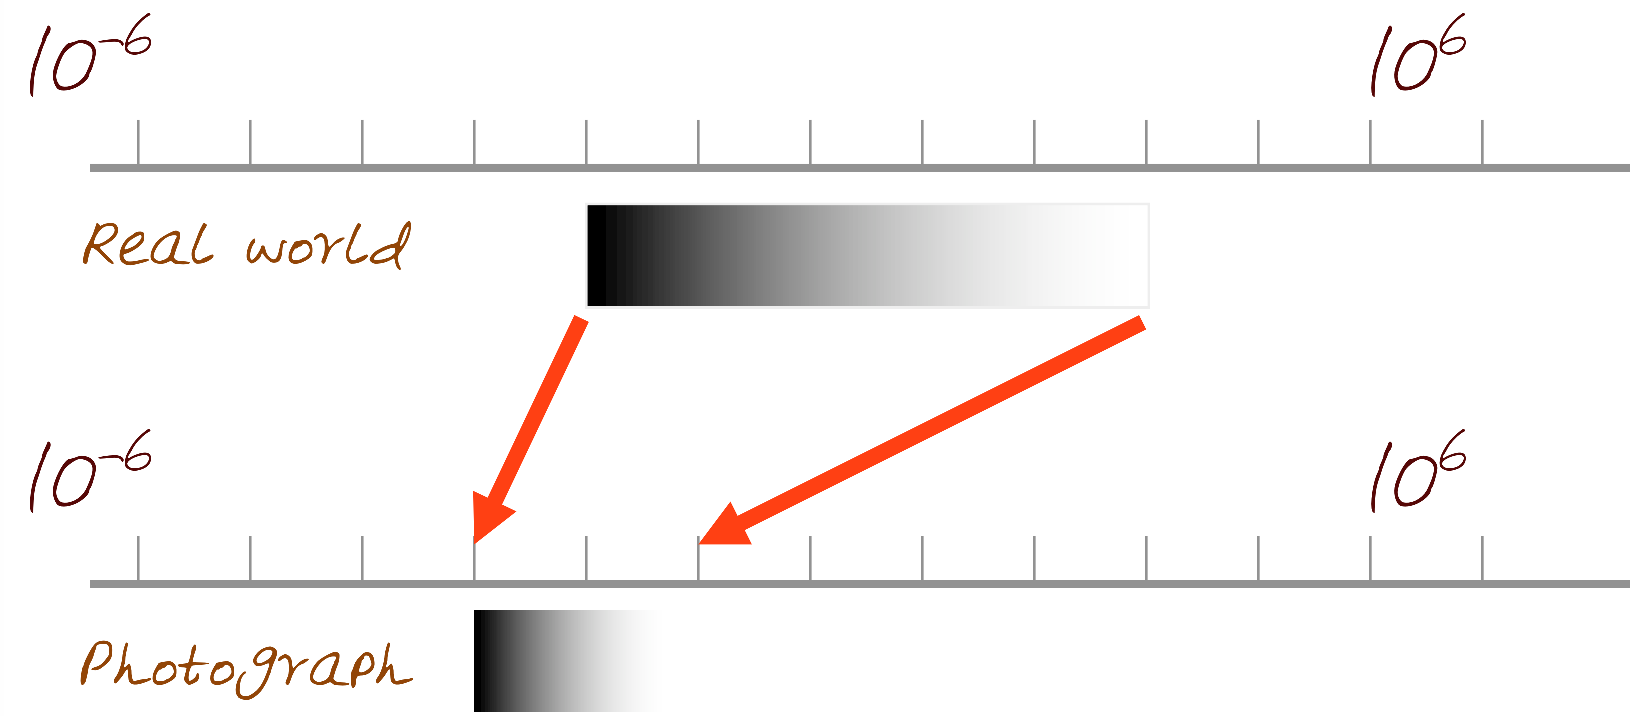

- 0 to 255

- Match limited contrast of the medium
- Preserve details

- Use filtering approaches to "compress" locally and globally

Further Information

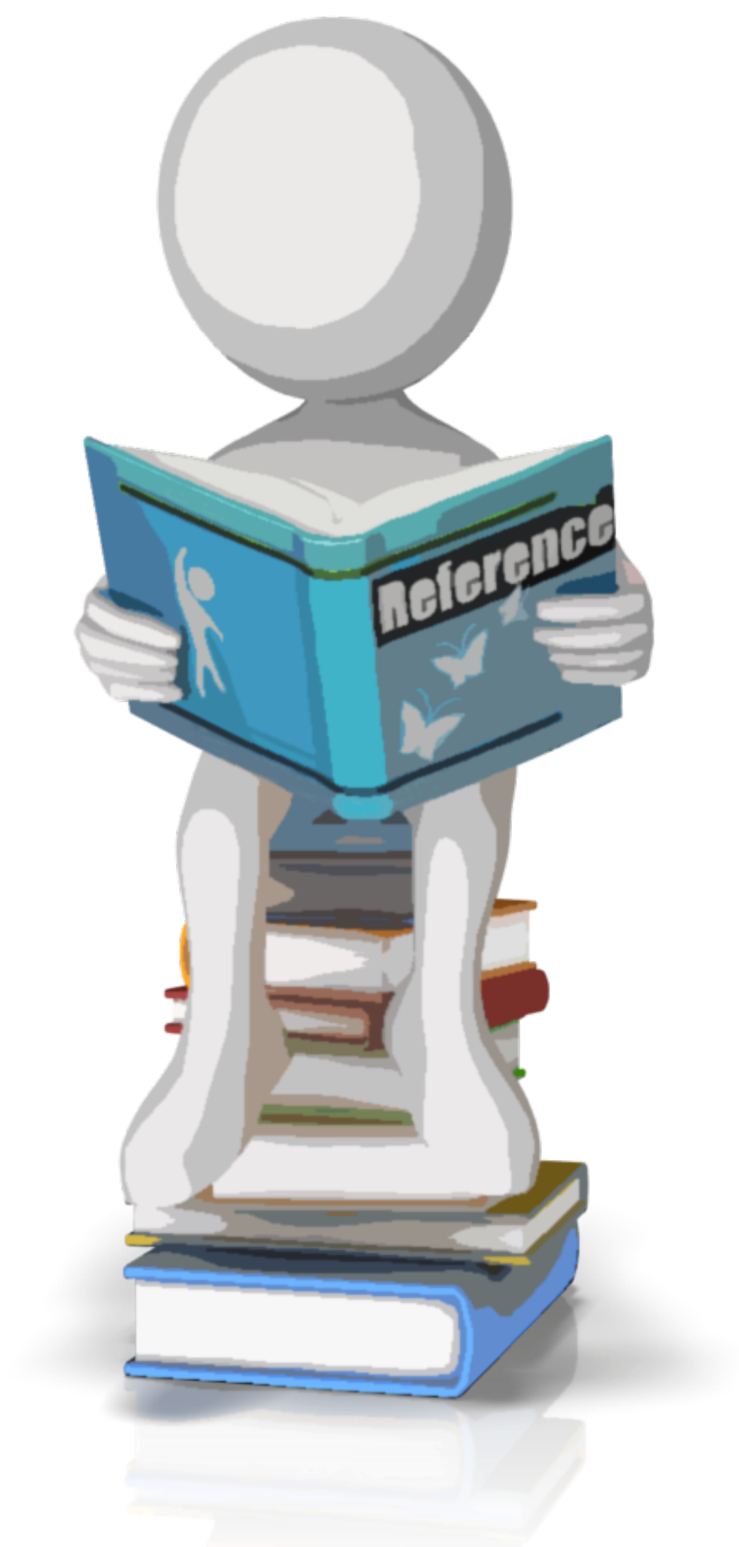

- Grossberg and Nayar (2003), "Determining the Camera Response from Images: What is Knowable?," IEEE Transactions on Pattern Analysis and machine Intelligence,
- Debevec and malik (1997). "Recovering High Dynamic Range Radiance maps from Photographs." In SIGGRAPH 1997
- Ward (2001), "High Dynamic Range Imaging," Proceedings of the Ninth Color Imaging Conference, November 2001.

Further Information

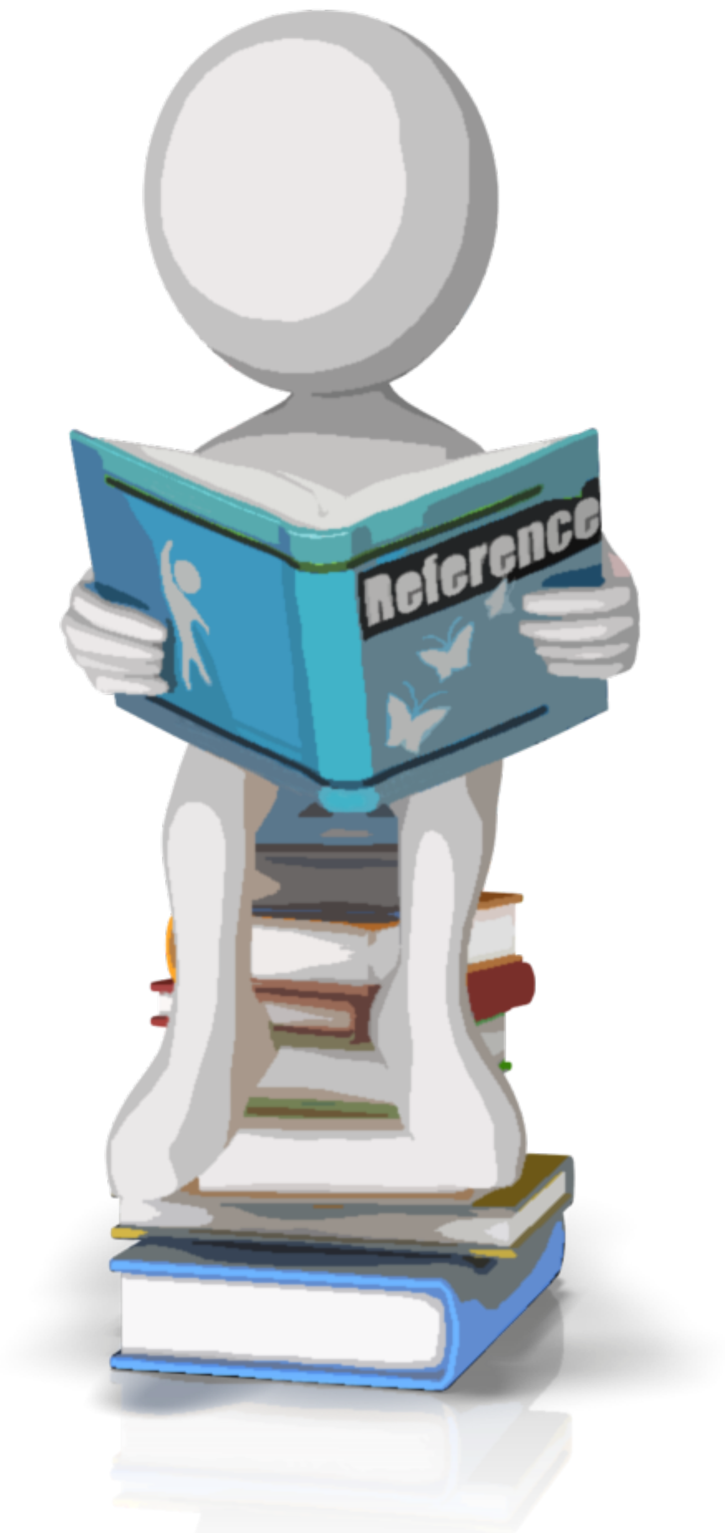

- Durand and Dorsey (2002), "Fast Bilateral Filtering for the Display of High-Dynamic-Range Images" In SIGGRAPH 2002.
- Reinhard, Stark, Shirley and Ferwerda (2002), "Photographic Tone Reproduction for Digital Images", In SIGGRAPH 2002.
- Banterle, Artusi, Debattista, and Chalmers (2011) Advanced High Dynamic Range Imaging CRC Press. (with matlab Code)
- many Software suites on the Internet.
- Also, look for "Exposure Fusion"

Credits

* Softwares used
  * matlab by mathwork's Inc.
* For more information, see
  * Richard Szeliski (2010) Computer Vision: Algorithms and Applications, Springer.
* Some concepts in slides motivated by similar slides by J. Hays.
* Photographs by Irfan Essa